



Universiteit
Leiden
The Netherlands

Frequency Stabilization of Four Cascaded Optical Cavities

Kanger, Hidde

Citation

Kanger, H. (2024). *Frequency Stabilization of Four Cascaded Optical Cavities*.

Version: Not Applicable (or Unknown)

License: [License to inclusion and publication of a Bachelor or Master Thesis, 2023](#)

Downloaded from: <https://hdl.handle.net/1887/3755816>

Note: To cite this publication please use the final published version (if applicable).



Frequency stabilization of four cascaded optical cavities

THESIS

submitted in partial fulfillment of the
requirements for the degree of

MASTER OF SCIENCE

in

PHYSICS

Author :	H. Kanger
Student ID :	s2332795
Supervisor :	Prof.dr. D. Bouwmeester
2 nd corrector :	Dr. W. Löffler

Leiden, The Netherlands, May 29, 2024

Frequency stabilization of four cascaded optical cavities

H. Kanger

Huygens-Kamerlingh Onnes Laboratory, Leiden University
P.O. Box 9500, 2300 RA Leiden, The Netherlands

May 29, 2024

Abstract

This research covers the development and application of multiple low-noise high-bandwidth lockboxes to position the mirrors of four cascaded optical cavities using piezo actuators for frequency locking with a precision of 40 Hz. The design and fabrication of a custom printed circuit board that hosts two ARM-based microcontrollers for signal processing are elucidated. By measuring the impedance of the piezo actuators over a large range of frequencies, we found several electromechanical resonances ranging from a few kHz to a main resonance at 80 – 100 kHz. It was found that these resonances greatly impact the ability to lock the cavity to the laser source and thereby impose bandwidth limitation on the feedback. By avoiding excitation of such resonances by reducing the feedback bandwidth to below the first prominent resonance at 2.5 kHz, we were able to achieve a high quality lock of a single optical cavity. Using a reduced bandwidth of 250 Hz and a reduced modulation frequency of 1.9 kHz, we demonstrated the locking of four cascaded cavities and achieved an optical transmission of $T \approx 40\%$, limited mostly by optical alignment. Finally, we show an initial lock freeze procedure for three cascaded cavities in which 90% of the transmission during the locked state is retained for a period of 4.9 seconds, while providing no feedback on the piezos.

Contents

1	Introduction	7
2	Experimental setup	11
2.1	Full experimental setup	11
2.2	Filter cavity setup	14
3	Theory	17
3.1	Dither locking	17
3.2	Pound-Drever-Hall	20
3.3	Z-transform	22
4	Lockbox hardware	25
4.1	Main board	25
4.1.1	Input amplification	27
4.1.2	Analog to digital conversion	29
4.1.3	Digital to analog conversion	30
4.1.4	Power	34
4.1.5	Grounding	34
4.2	I/O extension PCB	35
4.3	Manufacturing	37
5	Single filter cavity locking	39
5.1	Piezo-mirror system dynamics	39
5.2	Locking procedure	44
5.2.1	Experimental error signal	46
5.2.2	Locking	47
5.3	Feedback bandwidth limitations	51
		5

6	Four filter cavities locking	55
6.1	Locking procedure	55
6.2	Optical transmission	59
6.3	Lock freeze	59
7	Discussion	63
7.1	Recommendations	63
7.2	Conclusions	65
	Appendices	73
A	Cavity ringdown	75
B	Auxiliary figures	77

Introduction

Quantum mechanics and the theory of general relativity are both very successful modern theories of which incredible theoretical predictions like entanglement, black holes, gravitational waves and much more have been experimentally verified repeatedly [1–3]. Despite these achievements, these theories exhibit tensions in their compatibility, particularly in the description of massive superpositions, black holes and time [4, 5]. It is these incompatibilities that have inspired both theoretical and experimental efforts to test the interplay between quantum theory and gravity. As with all scientific theories, hypotheses emerging from these efforts must be falsifiable through experiment. However, the length scale difference between both theories and the relative weak nature of the gravitational force complicate experimental verification of existent theories. Nevertheless, various experiments have been proposed involving the superposition of massive objects to study the stability of such superpositions under gravitational interactions [6–8].

One such experiment is a membrane-in-the-middle setup where a 25 nm thick SiN membrane is put inside an optical cavity and cooled to mK temperatures [9]. In such a setup, vibrational modes of the membrane can be coupled to the field of the laser. Leveraging this optomechanical coupling, spatial superpositions of the membrane can be realized through a process called Stimulated Raman Adiabatic Passage (STIRAP) [9].

A fundamental effect underlying such an experiment is a photon-phonon interaction called Raman scattering. By interaction of a photon at frequency ω_p with a membrane phonon of frequency Ω_m , the incident photon is frequency shifted by the phonon frequency according to:

$$\omega_s = \omega_p \pm \Omega_m, \quad (1.1)$$

where ω_s is the scattered (anti-)Stokes photon that has either gained (anti-Stokes) or lost (Stokes) energy [10]. To realize STIRAP experimentally, it is essential to filter the photons of interest, the Stokes photons, from the so-called “pump light”. For the membrane studied in this research with $\Omega_m \approx 1$ MHz, an optical filter with an attenuation of 150 dB of the pump photons is required for efficient detection of the single (anti-)Stokes photons [11]. This would require a Fabry Perot cavity with a linewidth of less than 1 Hz. Experimentally, this is difficult to realize due to the quality of optical components and because this requires frequency control of the cavity to within a fraction of 1 Hz. Furthermore, a cavity with a linewidth of 1 Hz adds a delay to a pulse traversing the cavity of approximately 300 ms*. Such a delay is unacceptable for STIRAP experiments as the decoherence time of the membrane at 1 K is approximately 5 ms [9]. However, by cascading N cavities, this problem can be circumvented as the delay through such a system grows linearly in N while the attenuation grows exponentially in N .

To this end, a system of four cascaded optical cavities has been designed and built as described in the work of V. Fedoseev (2022) [9]. Each cavity has a measured linewidth of 22.44 kHz, yielding a total delay of approximately 60 μ s, which is within the design requirements. For any (anti-)Stokes photons to be transmitted through the filter cavity system, all four cavities should be kept on resonance with those photons. This can be accomplished by precisely controlling the length of the cavity using the piezo actuators that are mounted on the mirrors of the cavity [9].

Previous work by V. Fedoseev (2022) on a custom built lockbox has demonstrated its sufficient capabilities for successfully locking of multiple filter cavities [9]. However, this lockbox suffered from limited feedback bandwidth, electronic noise and overall ease of operation. This reduced the overall transmission of the filter cavity system to approximately 20%.

This work covers the development of an improved lockbox and the subsequent testing of the device on the filter cavity system. With an improved quality of the lock, a higher overall transmission can be achieved. By thorough consideration of electronic components, electrical grounding and shielding, a high-bandwidth low-noise lockbox is realized. Through impedance characterization of the piezo-mirror systems, insight is gained in the resonant frequencies of the system, which pose limitations on the bandwidth of the feedback loop. The locking of a single cavity demon-

*The reader is reminded of the quality factor: $Q = \frac{\omega_0 \tau}{2}$, where ω_0 the laser frequency and τ the ringdown time. The delay is then: $\tau \sim 1/\nu_{FWHM}$, where ν_{FWHM} the cavity linewidth. See appendix A for a more detailed treatment.

strates the sufficient capabilities of the lockbox and enables locking of consecutive filter cavities. By reducing the bandwidth of the feedback loop to prevent excitations of the aforementioned resonances, the locking of four cascaded optical cavities is demonstrated while achieving high optical transmission. Being able to lock consecutive cavities, a first lock freeze experiment is performed in which the transmission through three cavities is recorded in time while no active feedback is applied.

The experimental setup to carry out quantum optomechanical experiments as well as the filter cavity setup are covered in chapter 2. Theoretical subjects that are instrumental to successfully locking optical cavities, like dither locking, the Pound-Drever-Hall technique and the z -transform, are covered in chapter 3. The various components of the lockbox hardware and design choices are covered in chapter 4. Impedance characterizations of the piezo as evidence for physical resonances in the piezo-mirror systems, the realization of a first lock as well as bandwidth limitations imposed by aforementioned resonances are covered in chapter 5. Chapter 6 covers the locking procedure and transmission results of a successful lock of all four cavities. The results of a first lock freeze measurement are discussed. Finally, in chapter 7, we put forward recommendations and conclude this work.

Experimental setup

The first section of this chapter describes the full quantum optomechanics experimental setup which provides the context in which the filter cavities will be used. It describes a membrane-in-the-middle setup that can be cooled down to mK temperatures inside a cryostat. Having discussed the relevant background information, the second section focuses on the filter cavity setup and the lockboxes. The design and fabrication of the filter cavity system were primarily done by V. Fedoseev and H. van der Meer. Alignment and initial experimentation on the filter cavity system involved numerous contributors, notably X. Wei and L. D. Raabe.

2.1 Full experimental setup

The setup to perform quantum optomechanical experiments is depicted in figure 2.1. Inside the cryostat a cavity with a membrane-in-the-middle configuration is held at mK temperatures. A photograph of the cryostat can be seen in figure B.2 in appendix B. The membrane is sideband heated/cooled by a pump laser, which is the Mephisto Nd:YAG 1064 nm laser from Coherent. It has a spectral linewidth of less than 3 kHz [12]. The frequency of the pump laser ω_b is shifted up/down with respect to the resonant frequency of the cold cavity ω_c by two acousto-optical modulators (AOM) from Isomet, such that:

$$\omega_b = \omega_c \pm \Omega_m, \quad (2.1)$$

where $\Omega_m \approx 1$ MHz the frequency of the mechanical mode. This can be achieved by setting one AOM at frequency f and a second AOM at frequency $-f + \Omega_m$.

Resulting (anti-)Stokes photons are filtered from the pump laser by the filter cavity system with a designed attenuation of the pump laser light of more than 150 dB at 1 MHz spectral separation [9, 11]. Transmitted (anti-)Stokes photons are fed back into the crystal to be detected by a superconducting nanowire single photon detector (SNSPD).

The probe laser is the Mephisto S Nd:YAG 1064 nm laser from Coherent that has a linewidth in the kHz range [12]. It is held on resonance with a reference cavity that has a linewidth of 7.5 kHz by laser crystal piezo control provided by the commercial modular wideband laser controller lockbox from the Vrije Universiteit Amsterdam that uses the Pound-Drever-Hall method (PDH). The input to this lockbox is the PDH error signal created by phase modulation at 9.5 MHz provided by the Thorlabs EO-PM-NR-C2 electro-optic phase modulator (EOM). The demodulation of the reflection signal is performed by the Zurich Instruments HF2LI lock-in amplifier. The reference cavity is introduced because the frequency of the lasers used typically drift around on the order of MHz/min due to temperature fluctuations [12]. The filter cavities are held on resonance with the probe laser by four custom built lockboxes. The development of such lockboxes is the focus of this work. However, quantum protocols like STIRAP, as described by V. Fedoseev (2022), typically use laser pulses rather than continuous waves. Therefore, during the collection of Stokes photons, the lock is “frozen” as no continuous light is incident on the filter cavity system for the lockboxes to lock onto. During this frozen state the lockboxes should keep the cavities on resonance as accurately as possible while no reflection signal, and thus no error signal, is available for the lockboxes.

The various components are locked as follows: the probe laser is locked to the reference cavity by means of crystal piezo control using the PDH method, the cold cavity is locked to the probe laser using the PDH method controlling the length of the cold cavity with the mounted piezos, the big innolite is locked to the probe laser by inspection of the beating signal and the filter cavities are locked to the big innolite using PDH for the first cavity and dither locking for the consecutive three cavities using the mounted piezos for length control.




Figure 2.1: Schematic overview of the membrane-in-the-middle setup for quantum optomechanical experiments. The cryostat contains a cavity with a membrane-in-the-middle setup at mK temperatures [9]. The membrane depicted here is a simulation of the displacement of the membrane for a defect mode [9]. The pump laser is used to cool or heat the membrane through the Raman interaction for which the two AOMs together provide the required frequency shift Ω_m according to (2.1). One AOM is at frequency f , e.g. $f = 80$ MHz, the second AOM is at $-f \pm \Omega_m$, where $\Omega_m \approx 1$ MHz the frequency of the membrane mode. The path around the AOMs is utilized by the filter cavities for locking. The resulting single (anti-)Stokes photons are filtered from the pump laser by the filter cavity system and are fed back into the cryostat to be detected by the SNSPD. The shutter after the filter cavities prevent pump light of reaching the SNSPD once all cavities are locked. During the collection of single photons the shutter is opened. The probe laser is locked to the reference cavity using the Pound-Drever-Hall method for with the EOM provides the modulation. The reference cavity is a narrow-linewidth cavity with a linewidth of 7.5 kHz. It is used as a fixed reference frequency for the entire experiment: the probe laser is locked to the reference cavity, the cold cavity is locked to the probe laser, the pump is locked to the probe by inspection of the beating signal and the filter cavities are locked to the pump laser by the custom lockboxes.

2.2 Filter cavity setup

Figure 2.2 shows a schematic overview of the filter cavity system together with the custom built lockboxes used for locking the cavities. The system consists of four consecutive Fabry-Perot cavities inside a vacuum chamber, each with a measured linewidth of 22.44 kHz at 1064 nm. Each cavity has its own photodetector that detects the backreflection of the cavity. The first cavity has a PDA10CS-EC InGaAs amplified Thorlabs detector and is placed outside the vacuum chamber. The other three detectors are inside the vacuum chamber and are SM05PD4A mounted InGaAs photodiodes that are unamplified and therefore require a transimpedance amplifier (TIA). This amplifier converts the small output currents of the detector to a reasonable voltage in the order of volts such that the signals can be used by the lockboxes. The TIA used is the Thorlabs AMP120 with a gain of 100 kV/A.

The individual cavities are decoupled by using free space optical circulators consisting of two $\lambda/4$ waveplates and a polarizing beam splitter. This removes backreflections of consecutive cavities on the detectors, simplifying the locking process significantly [9, 11]. The high quality mirrors have low absorbance and scattering of < 3 ppm. The designed transmission through the polarizing beam splitters (PBS) and waveplates in the optical isolators is 96% [9]. These considerations are essential to achieve high transmission through the entire system that contains many optical surfaces.

The filter cavity system is approximately 2 m long, a photograph can be seen in figure B.1 in appendix B. Each cavity has a length of 40 cm. The measured reflection signal of each detector is fed into its own lockbox that calculates the required PID feedback and outputs the right voltage on a piezo attached to the input or output mirror of the cavity to keep the cavity length on resonance. Both mirrors of the cavities have a piezo mounted to it for control. This is mainly used for intra-cavity spacing but due to the free-space circulators the cavities are effectively decoupled, removing the dependency of the transmission on the intra-cavity spacing [11]. The piezos used are the NAC2124 ring piezo actuators with a maximum free stroke of $3.3 \mu\text{m}$ at a maximum voltage of 200 V. The transmitted Stokes photons are coupled into a fiber and fed back into the cryostat to be detected by a superconducting nanowire single photon detector (SNSPD). The shutter after the last cavity will only be in the open state during the collection of single photons as to protect the SNSPD from the strong pump light. During this period, the locks are frozen as no pump light is incident on the filter cavities.

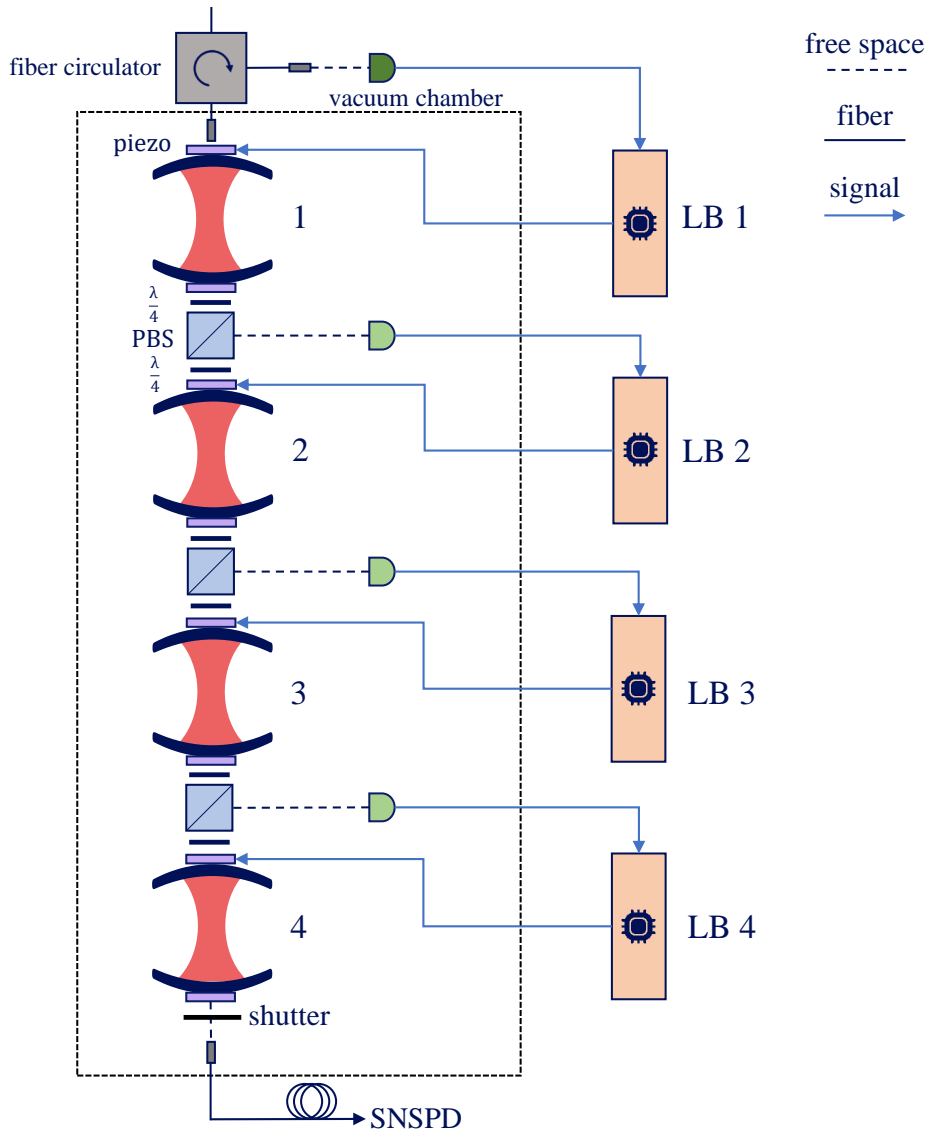


Figure 2.2: Schematic overview of the experimental setup of the filter cavity system. The 2 meter long vacuum chamber hosts four consecutive optical cavities with mirror-piezo mounts for precise control of the cavity length. Each cavity has a length of 40 cm. The cavities are decoupled from each other using free space optical isolators such that the signal incident on the detector only contains the reflection signal of a single cavity. The light inside the cavities is circularly polarized. Each cavity has its own lockbox that takes as input the reflected signal for that cavity and outputs the voltage on the piezo attached to the input or output mirror of that cavity to keep the cavity at the right length. All connections to and from the vacuum chamber are fed through a single LEMO connection hub. A photograph of the system can be seen in figure B.1 in appendix B.

Theory

In order to understand the locking of cavities to a laser frequency, it is important to understand the theory behind locking techniques like dither locking and Pound-Drever-Hall. Furthermore, as digital computation is inherently discrete, the Z-transform is used to allow for implementation and numerical stability of Laplace transfer functions in the discrete domain. The mathematical backgrounds behind these concepts are treated in separate sections below.

3.1 Dither locking

The reflection \mathcal{R} of light on a single cavity as a function of the frequency ω is a Lorentz distribution of the form:

$$\mathcal{R}(\omega) = 1 - \frac{(1 - R)^2}{(1 - R)^2 + \frac{4RL^2}{c^2}(\omega_n - \omega)^2} \quad (3.1)$$

where $R = |r|^2$ the squared Fresnel coefficient of reflection of the mirrors, L the length of the cavity and ω_n are the resonant frequencies [13]. The resonant frequencies of the cavity can be derived from the condition of constructive interference inside a cavity of length L :

$$\omega_n = \frac{n\pi c}{L}, \quad (3.2)$$

where $n \in \mathbb{Z}$. Because it is the length of the cavity that is being feedback on, we rewrite (3.1) in terms of L :

$$\mathcal{R}(L) = 1 - \frac{(1 - R)^2}{(1 - R)^2 + \frac{4R\omega^2}{c^2} (L_n - L)^2}. \quad (3.3)$$

We see, from the $(L_n - L)^2$ term that it is a symmetric function around the resonance L_n . This means that if one were to try to provide active feedback to keep the length of the cavity on resonance by using the reflected signal only, one cannot know in “which way to steer” to correct for perturbations.

Instead, one would need the derivative of the reflection profile to be able to lock to the center of the resonance. This can be achieved by a technique called dither locking. By modulating the length of the cavity using a piezo on one of the mirrors at a frequency ω , you effectively modulate the reflection signal with the same frequency near the resonance:

$$R(t) = R_{res} + B \sin(\omega_m t + \phi), \quad (3.4)$$

where R the reflection, R_{res} the minimal reflection at a resonance, B the resulting dithering amplitude in the reflection signal, ω_m the modulation frequency and ϕ a phase factor.

Now, in order to retrieve the derivative of the reflection signal, the dither locking procedure involves demodulation as depicted schematically in figure 3.1. First, the incoming reflection signal, as described in (3.4), is multiplied with a local oscillator that has the same modulation frequency ω_m and a phase ψ :

$$\vec{X} = \begin{bmatrix} R(t) \cdot \sin(\omega_m t + \psi) \\ R(t) \cdot \cos(\omega_m t + \psi) \end{bmatrix} \quad (3.5)$$

$$= \begin{bmatrix} R_{res} \sin(\omega_m t + \phi) + \frac{B}{2} \cos(\phi - \psi) + \frac{B}{2} \cos(2\omega_m t + \phi + \psi) \\ R_{res} \cos(\omega_m t + \phi) + \frac{B}{2} \sin(\phi - \psi) + \frac{B}{2} \sin(2\omega_m t + \phi + \psi) \end{bmatrix}. \quad (3.6)$$

The resulting signal is passed through a low-pass filter such that the fast-changing components are filtered out and only the DC signal remains:

$$H_{LPF}(\vec{X}) \approx \frac{B}{2} \begin{bmatrix} \cos(\phi - \psi) \\ \sin(\phi - \psi) \end{bmatrix}. \quad (3.7)$$

We see that the resulting signal has two components that are $\frac{\pi}{2}$ out of phase, these are called the *in-phase* and *quadrature* components following the usual convention [14]. Because the phase difference $\phi - \psi$ is largely unknown and is setup specific, the vector is rotated as to maximize the in-phase component, indicated by the $Rot(\theta)$ block in figure 3.1.

Following O. Huisman (2021), we now show that this is indeed the quantity we are interested in [15]. We write:

$$L(t) = L_{res} + A \sin(\omega_m t + \phi), \quad (3.8)$$




Figure 3.1: Schematic overview of the dither locking procedure. The dashed lines represent optical paths while the solid lines indicate signal paths. The cavity is on the top right with a piezo attached to the right mirror. The laser comes in from the left. Following the path from the bottom left to right: the measured reflection signal is first mixed with a modulation signal, then passed through a low-pass filter after which the error signal is optimized by applying a rotation over an angle θ . The PID feedback is then fed back into the system together with the modulation signal.

where L is the length of the cavity, L_{res} a length where the cavity is on resonance, A is the modulation amplitude, ω_m the modulation frequency and ϕ the phase of the modulation. After expanding $R(L(t))$ around L_{res} one obtains:

$$B \approx A \frac{dR}{dL}. \quad (3.9)$$

Referring back to (3.7), we see that indeed the dither locking procedure successfully extracts the derivative of the reflection signal with respect to the length of the cavity. This signal is the error signal and serves as the input to the PID controller.




Figure 3.2: Schematic overview of the Pound-Drever-Hall procedure. The dashed lines represent optical paths while the solid lines indicate signal paths. The phase of the light is modulated using an Electro-Optic Modulator (EOM). Instead of providing feedback to a piezo on the cavity, the feedback is fed into the laser to stabilize the laser frequency.

3.2 Pound-Drever-Hall

A better performing and more popular technique is the so-called Pound-Drever-Hall technique (PDH hereafter), developed in 1983 by R.V. Pound, R. Drever and J. L. Hall [16]. This technique is often used to stabilize the frequency of a laser to a stable cavity to be used as a frequency reference [17]. Instead of modulating the length of the cavity, this technique modulates the phase of the light to generate an error signal. A schematic overview of the technique can be seen in figure 3.2. The laser frequency is stabilized by providing PID feedback on the piezo crystal control of the laser. The error signal is generated using phase modulation performed by the Thorlabs EO-PM-NR-C2 Electro-Optic Modulator (EOM). This device contains a birefringent lithium niobate crystal of which the refractive index along one of the axes is changed by applying an electric field. Providing the device with an AC voltage then results in phase modulation of the field. Following E. Black (2001), an insightful derivation for the Pound Drever Hall error signal is given.

The electric field of the laser light before and after the EOM can be described as:

$$E(t) = E_0 e^{i\omega t + \phi} \xrightarrow[\text{phase mod.}]{\text{EOM}} E_0 e^{i\omega t + \beta \sin \Omega t}, \quad (3.10)$$

where E_0 is the amplitude of the electric field, ω the frequency of the electric field, ϕ the phase of the electric field, β the modulation amplitude and Ω the modulation frequency. Using the identity:

$$e^{iz\cos\theta} = \sum_{n=-\infty}^{\infty} i^n J_n(z) e^{in\theta}, \quad (3.11)$$

known as the Jacobi-Anger expansion, one can write (3.10) as:

$$E(t) \approx E_0 \left(J_0(\beta) e^{i\omega t} + J_1(\beta) e^{i(\omega+\Omega)t} - J_1(\beta) e^{i(\omega-\Omega)t} \right), \quad (3.12)$$

where J_n the n -th order Bessel function [17]. The field, in the above approximation, then contains a carrier frequency at the original laser frequency and two sidebands shifted with $\pm\Omega^*$. When this field reflects of a cavity with transfer function $\mathcal{R}(\omega)$, as in (3.1), the reflected field is:

$$\begin{aligned} \mathcal{R}(\omega) \cdot E(\omega) = E_0 \left(\mathcal{R}(\omega) J_0(\beta) e^{i\omega t} + \mathcal{R}(\omega + \Omega) J_1(\beta) e^{i(\omega+\Omega)t} \right. \\ \left. - \mathcal{R}(\omega - \Omega) J_1(\beta) e^{i(\omega-\Omega)t} \right), \end{aligned} \quad (3.13)$$

where we made use of the superposition principle for the three electric fields. Then, after a lengthy derivation of the demodulation procedure that goes beyond the scope of this work, one arrives at an expression for the error signal:

$$\epsilon = -2\sqrt{P_c P_s} \text{Im} \{ \mathcal{R}(\omega) \mathcal{R}^*(\omega + \Omega) - \mathcal{R}^*(\omega) \mathcal{R}(\omega - \Omega) \} e^{i\theta}, \quad (3.14)$$

where P_c and P_s are the power of the carrier and sidebands respectively and $e^{i\theta}$ is an extra phase factor that serves as a parameter to optimize the amplitude of the carrier, similar to dither locking. It is interesting to note that according to (3.14), the (power of the) sidebands are crucial for the error signal. This is also the reason that this technique cannot be used to lock consecutive cavities because the sidebands are filtered out by the preceding cavity. Figure 3.3 shows an example of an error signal with modulation frequency ω_m and a resonance at ω_r . The error signal shows a clear linear region around the resonance of the cavity ω_r . This linear region is where the feedback loop should be initiated for the lock to be able to reach zero. Note that depending on the implementation of the PID controller, the slope of this linear region should be either positive or negative. Conforming to convention, we use a negative slope in this work.

*It is insightful to note that the amplitude of the carrier and the sidebands are described by Bessel functions of a different order, which are functions that have multiple zeros. This makes it possible to make the power in the carrier or sidebands completely disappear at certain modulation amplitudes. This is useful for controlling the power in the carrier and sidebands.

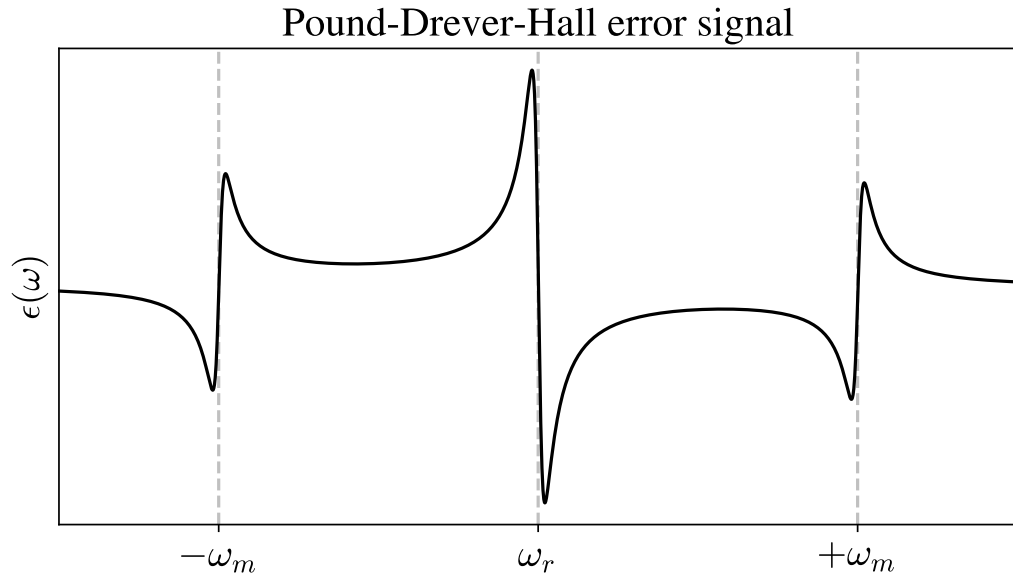


Figure 3.3: An example of a Pound-Drever-Hall error signal. The units are arbitrary. The small linear region in the center crosses zero at the resonance of the cavity ω_r , this is where the feedback happens. The two sidebands are clearly visible on either side and are at a distance ω_m . Note that depending on the phase of the demodulation, the slope of the Pound-Drever-Hall error signal can be positive or negative, which should match the definition of the PID implementation. Conforming to convention we use a negative slope in this work, as shown in the figure.

3.3 Z-transform

Because of the fundamental design of computers, the processing of signals is done in a discrete manner; signals are represented as a temporal sequence of numbers rather than a continuous variable as for analogue signals. As standard Fourier and Laplace transforms use continuous variables, it is important to give extra attention to the processing of discrete signals.

In the Laplace domain one has a simple relation between the input signal $X(s)$ and output signal $Y(s)$ of a linear time-invariant system (LTI):

$$Y(s) = H_c(s)X(s), \quad (3.15)$$

where $H_c(s)$ is the transfer function of the system in the Laplace domain. However, for a discrete system, we are interested in a so-called “difference

equation". Such an equation has the form:

$$\sum_{k=0}^N a_k y[n-k] = \sum_{l=0}^M b_l x[n-l], \quad (3.16)$$

where $x[n]$ and $y[n]$ are the system input and output signals at discrete time step n respectively and a_k and b_l the coefficients of these terms. The attentive reader might notice that this is also a general form of an infinite impulse filter where the output depends on all previous in- and outputs. One can rearrange (3.16) to obtain a perhaps more insightful form:

$$y[n] = \sum_{l=0}^M b_l x[n-l] - \sum_{k=1}^N a_k y[n-k]. \quad (3.17)$$

To arrive at such an equation one can use the z-transform [18]. The z-transform of a discrete signal $x[n]$ is defined as:

$$X(z) = \mathcal{Z} \{x[n]\} = \sum_{n=-\infty}^{\infty} x[n]z^{-n}, \quad (3.18)$$

where $n \in \mathbb{N}$ is the index of the discrete time signal and z is a complex number. Before applying the (inverse) z-transform, the transfer function has to be transformed to a discrete time transfer function using the bilinear transformation. The exact mapping from the s -plane to the z plane is defined as:

$$s \longrightarrow \frac{1}{T} \ln z, \quad (3.19)$$

where s is the s -plane variable of the Laplace domain, z the z -plane variable of the z domain as in the z-transform and T_s the sampling period of the discrete system [18]. The bilinear transform uses the first order approximation of this mapping:

$$s \approx \frac{2}{T_s} \frac{z-1}{z+1}. \quad (3.20)$$

The transfer function is then transformed as:

$$H_d(z) = H_c \left(\frac{2}{T_s} \frac{z-1}{z+1} \right), \quad (3.21)$$

where $H_d(z)$ is the discrete time transfer function. After evaluation of (3.21) for an LTI system, one can gather all terms and use:

$$H_d(z) = \frac{Y(z)}{X(z)} = \frac{b_0 + b_1 z^{-1} + \dots + b_M z^{-M}}{a_0 + a_1 z^{-1} + \dots + a_N z^{-N}}, \quad (3.22)$$

where $X(z)$ and $Y(z)$ are the input and output signals in the z -plane respectively. Rearranging those terms yields:

$$Y(z) = \left(b_0 + b_1z^{-1} + \dots + b_Mz^{-M} \right) X(z) - \left(a_1z^{-1} + \dots + a_Nz^{-N} \right) Y(z). \quad (3.23)$$

Now, one can identify the z -transform and arrive at an equation of the desired form:

$$y[n] = b_0x[n] + b_1x[n-1] + \dots + b_Mx[n-M] - a_1y[n-1] - \dots - a_Ny[n-N]. \quad (3.24)$$

In this form the transfer function can be implemented digitally on e.g. a micro-controller. One might argue that for a simple PID controller the above method is disproportionate as one can also simply implement the integral and derivative terms naively by using Eulers method. However, the Z -transform assures numerical stability when going from the continuous Laplace domain to a discrete system. The Euler method does not provide such guarantees. Furthermore, the above method can be applied to any general transfer function where a naive implementation is not obvious, e.g. filters.

Chapter 4

Lockbox hardware

This chapter covers the printed circuit board (PCB) of the lockbox and the design choices made to arrive at the design presented in this work. The first section covers the main features of the board. The second section covers an extension PCB developed to allow for input and output to and from the Teensys on the main board. Lastly, the manufacturing process is described.

Much of the designs for both the hardware of the main board and software for a first version were previously developed by E. Morell. Although the original concept, as described by V. Fedoseev (2021), served as inspiration, the content, structure and execution of this lockbox has been fully reimaged by E. Morell.

4.1 Main board

The main PCB design is covered in this section, of which the most important parts are in their own sections covering the input, analog to digital conversion, digital to analog conversion and how the power is managed on the board. Extra attention is given to the grounding scheme throughout the PCB in a separate section at the end. The board hosts a mix of surface mount devices (SMD), integrated circuits (IC), connectors and more, amounting to a total of more than 220 individual components. For a general overview of the PCB board, see figure 4.1 below. Two Teensys are mounted on the bottom of the PCB with through-hole pins as can be seen in the top left of figure 4.1. The input reflection signal comes in at the white SMA connector at the bottom of the input section. After amplification, the signal is fed to both Teensys. All signal processing and digital communi-

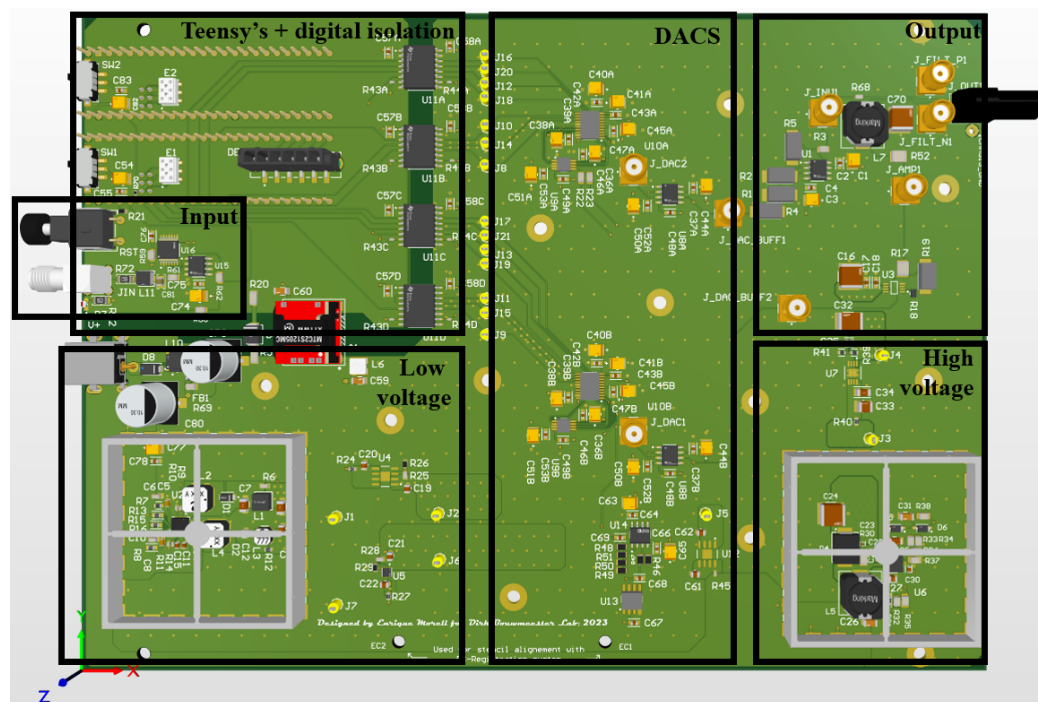


Figure 4.1: Annotated 3D render of the main PCB board. The white SMA connector in the bottom left of the input section receives the reflection signal and is amplified accordingly, as discussed in section 4.1.1. The board is powered through the 1.75 mm jack connector in the top left corner of the power section. Two Teensys are mounted on the bottom of the PCB with through-hole pins, as can be seen in the top left of the Teensy section. Two DACs are digitally addressed through the isolators and subsequently amplified and combined in the LEMO output connector that connects to the piezo.

cation with the DACs is performed by the bottom Teensy while the second Teensy can offload tasks from the main Teensy, e.g. host an HTTP-server allowing for communication over ethernet with the Teensys for sending live signals, status updates, configurations and more. The two white female headers (E1 & E2) on both Teensys and the horizontal black header on the bottom Teensy allow for ethernet and general signal input/output connections through a separate PCB, which is described in more detail in section 4.2. After digital processing of the signal, the calculated output is communicated with the DACs through the digital isolators, separating the ground planes. The output of both DACs are then adequately amplified and combined in the LEMO output connector that connects to the piezo.

4.1.1 Input amplification

Because the input amplitude of the incoming reflection signal is not of equal amplitude for all cavities or from run to run due to external factors, it is important to normalize the signal that comes in to a preset level such that important parameters in the system, e.g. PID parameters, are independent of the incoming signal strength. Furthermore, to optimize the accuracy of the ADC, it is best to use the entire input range of the ADC for the incoming signal. For these reasons an input amplification circuit is designed. At startup of the lockbox the incoming signal will be amplified to the desired level. A render of the input section of the PCB can be seen in figure 4.3.

The amplification circuit uses the OPAx388 precision opamp in a closed-loop amplifier configuration where the resistances are controlled by the dual channel AD5252 digital potentiometer as in figure 4.2. The AD5252 features two digitally addressable wipers that determine the resistance on both channels. The wipers are controlled by writing data to the volatile RDAC registers that can be stored to the non-volatile EEMEM registers such that they are reloaded on power-up when desired. Software has been written to communicate between the main Teensy and the digipot over the I2C protocol, a synchronous master-slave communication protocol typically used for communication between integrated circuits at low speeds. This is a suitable protocol because communication only happens during calibration which typically occurs only once per run.

By controlling the feedback resistor R_f and the ground resistor R_g , the amplification created is then:

$$\frac{V_{out}}{V_{in}} = 1 + \frac{R_f}{R_g}, \quad (4.1)$$




Figure 4.2: Circuit diagram for the input amplification. The amplification is controlled by the dual channel AD5252 digital potentiometer controlling the resistances of the feedback resistor R_f and the ground resistor R_g . The amplification of this circuit is then given by $V_{out}/V_{in} = 1 + R_f/R_g$.

where V_{in} and V_{out} are the input- and output voltages respectively.

Because the Teensy 4.1 features an ADC with a resolution of 12 bits, the maximum value, in arbitrary units, it can return is $2^{12} = 4096$. Since it has been observed that amplification (close) to this limit increases noise and to avoid saturation of the ADC at which point the calibration is not accurate anymore, a desired input amplitude of $A_{desired} = 3500$ is used. After averaging the measured incoming signal amplitude over a few seconds on startup, the following optimization problem is solved by the Teensy:

$$\begin{aligned} & \text{minimize} && \left| A_{desired} - \left(1 + \frac{F_{\text{bit} \rightarrow \text{resistance}}(B_f)}{F_{\text{bit} \rightarrow \text{resistance}}(B_g)} \right) \right| \\ & \text{subject to} && 0 \leq B_f \leq 255 \\ & && 1 \leq B_g \leq 255, \end{aligned} \quad (4.2)$$

where B_i is the bit code representation of the resistance for the ground and feedback resistors and

$$F_{\text{bit} \rightarrow \text{resistance}}(B) = \frac{B}{256} \cdot 50\text{k}\Omega + 75\Omega, \quad (4.3)$$

a linear function inherent to the AD5252 digipot that relates the wiper bit code to the physical resistance on that channel [19]. For the reason that the calibration procedure is typically ran only once during initialization of the lockbox, the optimization problem is solved by calculating all combinations of resistance values and picking the combination that results in an amplification that is closest to $A_{desired}$.

Although the wipers can only take on 256 different values, simulations show that this method gives accurate amplifications to within 0.05% of the desired amplification for over 95% of amplifications between 1 and 20, the worst being only 0.175% off. This is more than sufficient for this application.

4.1.2 Analog to digital conversion

After input amplification, the signal is digitized by the 12-bit analog-to-digital converter (ADC) on the Teensy 4.1. To prevent aliasing, a second-order low-pass filter is placed after the input. More specifically, it is an RLC filter with $R = 50 \Omega$, $L = 12 \mu\text{H}$, and $C = 0.018 \mu\text{F}$, resulting in a cut-off (-3 dB) at approximately 230 kHz. This limits the ADC sampling speed to a minimum of 460 kHz by the Nyquist theorem [20]. The highest setting for the ADC sampling results in a measured sampling frequency of $f_s \approx 1.33 \text{ MHz}$, which is far above the minimum sampling speed required.

The sampling is performed using a technique called Direct Memory Access (DMA). Usually, the central processing unit (CPU) on the Teensy does one task at the time and is unavailable during computation. However, using DMA, data is buffered in memory and only sent to the CPU once an interrupt is raised indicating that the buffer is full. This frees up the CPU to perform other tasks while sampling data, effectively allowing the Teensy to simultaneously sample and process data. However, this does impose a requirement on the processing time for a buffer as it should be smaller than the time in which the next buffer is sampled. In practice, the requirement is the other way around where, assuming a realistic lower bound on the processing time of a single buffer, the buffer size should be large enough in order to allow for the CPU to process the previous buffer. For example, it was found that if an HTTP server was hosted on the main Teensy, a buffer size of 16 data points was found to be the minimum.




Figure 4.3: Render of the input stage of the lockbox. The white SMA connector receives the input reflection signal. After passing through the RLC filter the signal is amplified by the circuit as in figure 4.2. The digipot (U16) has two channels that adjust the resistances used by the opamp (U15) to amplify the signal. The black button (RST) serves as a general purpose physical input to the main Teensy.

4.1.3 Digital to analog conversion

The analog output of the lockbox is generated using two digital to analog converters (DAC). To this end, two AD5791 high performance 20 bit DACs are used. The AD5791 has a integral nonlinearity (INL) of ± 1 least significant bit, indicating that any (non-linear) output deviations from the theoretical value are within one least significant bit [21]. Furthermore, the slew-rate of the output is $50 \text{ V} / \mu\text{s}$, which quantifies the speed at which the DAC is able to change the output voltage and shows that it is suitable for high-speed operation.

The board hosts two DACs: DAC 1 takes a positive reference voltage of 5 V and DAC 2 takes a positive reference voltage of 10 V. The output of DAC 2 gets amplified to 50 V by a circuit similar to the input amplifier as in figure 4.2 but instead with fixed resistances with a ratio of four. The output of DAC 1 is inverted to -5 V using a standard inverting opamp configuration. The output voltage of DAC 1 is negative because the output voltages of both DACs are combined in the negative and positive pin respectively

of a LEMO connector which goes to the piezo. This piezo actuator should never receive a negative potential as this could break the device. By making the voltage of DAC 1 negative, it is made sure that the polarity of the pins is never reversed which could otherwise easily happen if the output of DAC 2 were to be higher than DAC 1 at any point during its operation.

Given the supply voltages on both DACs, the output resolution for DAC 1 and DAC 2 are approximately $5\ \mu\text{V}$ and $50\ \mu\text{V}$ respectively. The free spectral range (FSR) of a single filter cavity is approximately 375 MHz. Given that the linewidth of a single filter cavity is 22.44 kHz, corresponding to approximately 3 mV applied on the piezo, one can scan an entire FSR with approximately 50 V while applying feedback at a resolution much higher than the linewidth of the cavity.

Figure 4.4 shows a render of part of the PCB for DAC 1.




Figure 4.4: Render of part of the PCB board that contains digital isolators and DAC 1. The digital isolators (U11A-B) are on the left and DAC 1 (U10A) is the larger IC in the center. The yellow points with the J designators are test points. The buffer unity-gain amplifier uses OPA182IDT and is located at the bottom right (U8A). The capacitors in this figure are mostly decoupling capacitors. An identical circuit is designed for DAC 2 and used the two other digital isolators (not in this render).

As to isolate the ground of the Teensy and input signal from the rest of the board while still allowing communication across the grounds, two quad-channel ISO7841FDWR digital isolators are used per DAC. The digital isolators are small transformers that capacitively couple signals from

one side to the other and can therefore have separate grounds on either side while still allowing communication [22].

The capacitors on the top pins of the isolators and around the DAC are decoupling capacitors that compensate for sudden current draws from the IC or elsewhere on the board. They do this by providing temporary current and compensation using stored charges. This is desirable because generally the traces from the supply to the IC are long and thus have relatively high inductance, reducing the ability to draw sudden currents. The larger yellow box-like capacitors (e.g. C40A, C41A, etc.) are 10 μF tantalum capacitors providing decoupling for low frequency while the smaller capacitors (e.g. C43A, C47A, etc.) are 0.1 μF ceramic capacitors, which are used for high frequency decoupling [23, 24]. This provides the circuits with a more stable and decoupled power supply.

The AD5791 is an unbuffered DAC which would make the output voltage load dependent because of the high output impedance of the DAC [21]. To this end, the output of the DAC is buffered with the highly dynamic OPA182IDT precision opamp. This opamp can provide sudden currents that could be drawn by the load, resulting in an accurate output voltage [25].

Communication with the DAC happens through the channels of the digital isolator at a clock speed of 35 MHz through the Serial Peripheral Interface (SPI) protocol. Because the clock signal frequency is high, an isolation technique called “via-fencing” is employed where vias are placed along the signal trace in close proximity as to form a barrier for electromagnetic radiation. Vias are holes in the board connecting PCB layers vertically. In this case, the top plane is connected to the ground plane to provide short return paths. The via fence can be observed on the trace passing through J14 in figure 4.4.

Together, these considerations result in a quiet and accurate DAC output. A spectrum of the output of DAC 1 for the relevant frequencies up to 100 kHz while modulating at $f_m \approx 8.25$ kHz with an amplitude of 1 mV has been plotted in figure 4.5.

The spectrum shows a strong peak at f_m and multiple harmonics at multiples of f_m . Even though a perfect sine wave would not have harmonics, the strongest harmonic is the first harmonic which has an approximate attenuation of 22 dB compared to the carrier. These are caused by the imperfect sinusoidal waveform generated by the DAC that inherently uses discrete steps in its output. Naturally, higher modulation amplitudes suffer less from this because of the minimum step of 5 μV . The noise floor during modulation is at approximately 3 μV while the noise floor without modulation is at 2 μV . The slight increase in noise during modulation




Figure 4.5: Spectrum of lockbox output during a 1 mV sinusoidal modulation at 8.25 kHz. A large peak at the modulation frequency is visible with smaller harmonics at multiples of f_m , with the strongest harmonic being the first harmonic with an approximate attenuation of 22 dB compared to the carrier. The noise floor during modulation is at approximately $3 \mu\text{V}$ while it is at approximately $2 \mu\text{V}$ without modulation. This is below the least significant bit output voltage of $5 \mu\text{V}$. The inset plot shows the spectrum up to 100 Hz where no strong peaks at 50 Hz are visible. Note, a voltage offset of -2.5 V for DAC 1 during this measurement causes large amplitudes for frequencies near zero.

is also attributed to the imperfect representations of the sinusoidal waveform causing other frequencies than f_m to be present in the signal. Nevertheless, both of the noise floors are well below the least significant bit output of DAC 1 ($5 \mu\text{V}$). Furthermore, the inset plot shows that there are no strong peaks at 50 Hz, indicating that there are no grounding problems in the setup and that no 50 Hz radiation is picked up by the board or cables. Note, the spectrum is measured while the output of the lockbox is connected to the SR770 FFT network analyzer while in reality it will be connected to the filter cavity system during experiment. Nevertheless, the connection is made such that in both cases the output is grounded to the building ground through the shielding of the output cable. Furthermore, no 50 Hz noise has been observed while performing measurements, indicating that indeed the grounding is similar.

4.1.4 Power

The entire board is powered using the 1.75 mm power jack at 15 V by a standard floating power supply, e.g. the ES150 from Delta Elektronika. When operated, the entire system draws approximately 200 - 300 mA, resulting in a power consumption of 3 – 4 W. After the power supply is filtered by a common-mode choke, the switching power supply in the housing in the bottom left of the PCB generates a stable power supply for the rest of the board at 15 V. It utilizes the ADP5070 switching regulator and is largely based on the EVAL-AD5791SDZ evaluation board that contains a power solution for the AD4791 DAC [26]. Two low-dropout regulators are used to reduce switching noise, resulting in ± 14 V. The ± 14 V is used for the positive and negative analog supplies of the DACs as well as the inverting opamp at the output of DAC 1. A separate high voltage switching power supply is in the housing in the bottom right of the PCB that supplies 60 V for the amplifying circuit of DAC 2. The housings on the power supplies are there to shield the rest of the electronics from electromagnetic radiation coming from the switching in the MHz regime. A final low-dropout regulator is used to convert the 60 V from the high voltage supply to 55 V, which is the maximum supply voltage for the ADA4522 opamp used in the amplifier circuit for DAC 2 [27].

4.1.5 Grounding

The ground plane of the board is separated in two sections: the input plane and the output plane. The input plane contains the Teensys and the input reflection signal, this is called the Teensy ground (TGND). The output plane contains the DACs, power supplies and amplifiers, this is called the analog ground (AGND). See figure 4.6 for a schematic overview of the grounding. The two separate planes can also be seen in figure 4.1 where the dark green region separates the two lighter green sections.

Digital isolators prevent currents from flowing through the PCB while still allowing communication across the two planes. The SMBJ33CA-TR transient-voltage-suppression (TVS) diode protects the two ground planes from diverging beyond the breakdown voltage of 36.7 V, which could happen if one of the planes would get charged over time. Eventually, both planes are connected to a common building ground: TGND through the detectors and input cables and AGND through the shielding of the LEMO cable that is connected to the metallic filter cavity vacuum chamber. Furthermore, the analog ground plane is connected to the housing of the lockbox by the mounting holes that have the ground plane exposed, as can be

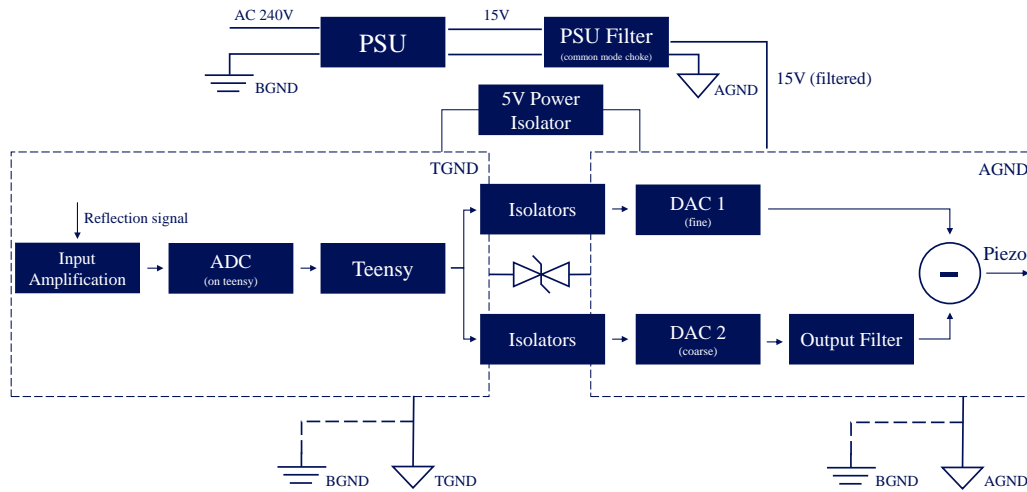


Figure 4.6: Schematic overview of the grounding scheme of the PCB. The board contains two planes: the Teensy ground (TGND) and the analog ground (AGND). The board is powered through a single power jack at 15 V on the analog plane. The Teensy is powered through the 5 V power isolator. To allow for communication across the planes but disallow currents to flow through the board, digital isolators are placed between the two planes. A TVS diode with a breakdown voltage of 36.7 V keeps the potential difference of the two planes within bounds. Through the detectors and the optical table, TGND is connected to the building ground (BGND). Through the shielding of the LEMO output cable and the metallic filter cavity vacuum chamber, AGND is also connected to BGND.

seen in figure 4.1. The small grid-patterned vias across the entire analog ground plane are part of a via-stitching, providing a short return path for currents which reduces impedances in the ground plane, resulting in an even ground plane throughout the entire PCB.

4.2 I/O extension PCB

To provide general input and output to and from both Teensys, an I/O extension board has been developed. The board contains two ethernet ports for communication over the internet and five SMA ports for general communication to and from the Teensy, e.g. input of an externally created error signal for the Pound-Drever-Hall method or triggering the oscilloscope based on events in the software. The extension board is connected to the two Teensys on the main board using custom built ribbon cables.

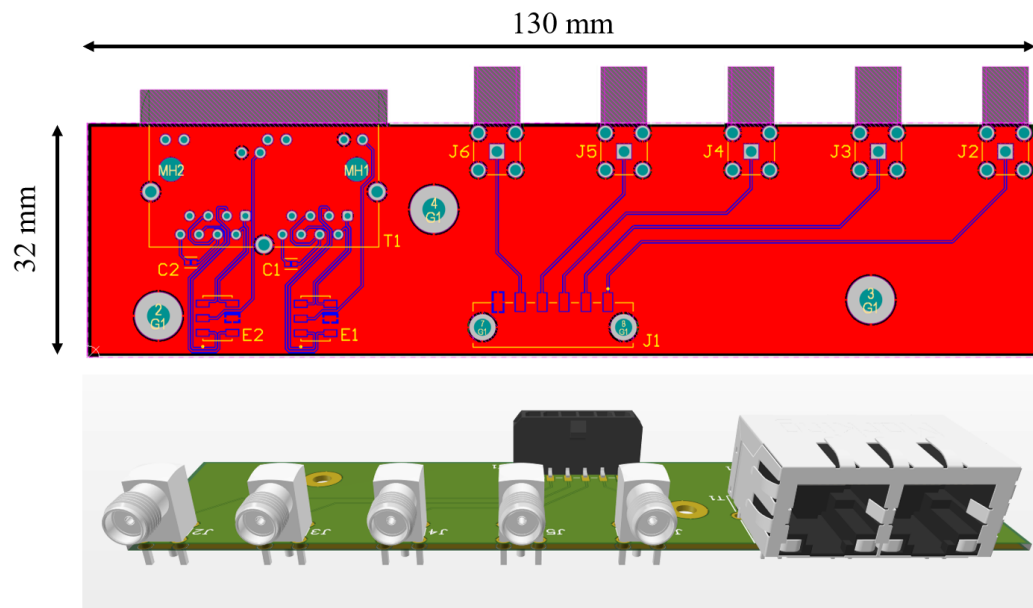


Figure 4.7: Render of the I/O extension PCB design. The board supports two ethernet connections and five SMA connections. The ethernet ports (T1) are connected to two low profile surface mount female headers (E1 and E2). The SMA connectors (J2 - J6) are connected to a six pin male SMT header. Using custom built cables these headers are connected to identical headers hosted on the main PCB. The red color in the top render is the top layer of the board, the blue/purple color around the traces is the ground plane of the board. The three large holes are mounting holes that also connect to the ground plane.

The ethernet connector is based on the Teensy 4.1 ethernet kit by PJRC*. See figure 4.7 for renders of the PCB design.

The board has a shared ground plane throughout the entire PCB, connecting directly to the Teensys through both the (internal) ethernet cable and the cable for the SMA ports. This ensures that the signals are referenced to the same ground, reducing mains hum. The ground of the extension board is also connected to the main board through the three mounting holes that have the ground plane exposed.

Conforming to the IEEE 802.3 standard for the physical layer of ethernet, the six pins on the ethernet connector contains a pair for transmission (TX+, TX-), a pair for receiving (RX+, RX-), a ground pin and an LED pin [28]. The transmission and receiving pairs are both differential pairs that are connected to a twisted pair of wires inside a standard ethernet cable through the ethernet header on the board. Unlike a reference to ground

*See https://www.pjrc.com/store/ethernet_kit.html for more information.

that is required for many signals, a differential signal carries the signal on both wires but at opposite potentials, e.g. -5 V and +5 V. This makes the communication less susceptible to common-mode noise as a simple subtraction of the two signals retrieves the original signal back, as depicted in figure 4.8.




Figure 4.8: Schematic overview of common-mode noise filtering. Using an opamp or common-mode choke, depicted by the triangle, the two signals can be subtracted to retrieve the original unperturbed signal back. It is important that the noise on the two wires in the differential pair is in phase for the subtraction to retrieve a clean original signal. This can be achieved by running the two traces close to each other or twisting the pair together when possible.

As evident from figure 4.8, for the subtraction of the two signals to cancel the noise optimally, it is important that the noise on the two wires in the differential pair is in phase. To this end, the data traces of the transmission and receiving pins are in close proximity and length-matched to within 1 mm, as can be seen in the bottom left of the top render in figure 4.7.

4.3 Manufacturing

Both the main- and the extension PCBs are manufactured externally. The board is assembled in-house and involves a few different steps.

Firstly, an SMD-stencil is placed on the PCB such that only the SMD pads are exposed. Leaded solder paste from Loctite is applied to the SMD pads by scraping it over the stencil such that only a thin layer of solder paste is left on each pad. Then, after careful removal of the stencil, the SMD components are placed on the board using the Ersas IR 550 rework station that allows for precision placement for small components. Next,

the entire PCB is placed in a reflow oven that uses vaporized Galden LS230 liquid, specifically for vapor phase soldering processes. It results in an even heat distribution for the entire PCB and allows for oxygen free soldering. It reaches a temperature of around 230 °C at which point the solder paste melts and flows freely, connecting the pads to adjacent metal surfaces. After cooldown, all connections are carefully inspected for shorts and fixed if necessary by either manual soldering or using a local reflow using the rework station. If all is well, the through-hole components are manually soldered. Lastly, the custom-built cables for ethernet and input/output through the extension board and the custom power cable are assembled. The final assembled lockbox together with the I/O extension PCB can be seen in figure 4.9.




Figure 4.9: Photographs of the lockbox and I/O extension board. The assembled main PCB (A), the assembled I/O extension PCB (B), the front-panel of the assembled lockbox (C) and the inside of the lockbox with the I/O extension PCB installed (D).

Single filter cavity locking

This chapter covers the analysis of non-linear regions in the piezo response caused by electromechanical resonances. Impedance measurements are compared to the reflection signal during the locking of the cavity to highlight the link between peaks in the impedance curves and physical resonances of the piezo-mirror system. The procedure to lock a single cavity to the probe laser is discussed in the second section. Finally, important bandwidth limitations caused by the aforementioned resonances are discussed.

5.1 Piezo-mirror system dynamics

A PID controller works under the assumption of a linear time-invariant system. Therefore, it is important that the piezo expansion reacts linearly to the applied voltage. However, it is known that piezos show electromechanical resonances [29]. These can be understood by considering the coupling of the displacement of a mass to the strength of the electrical field inside the piezoelectric material. This sets the scene for a typical resonant system. Resonances are highly non-linear responses which must therefore be avoided during the feedback process.

Each mirror of the cavity has a Noliac NAC2124 ring piezo actuator attached to it, which are reported to have an unloaded resonance at approximately 486 kHz [30]. However, because a mirror is mounted to the piezo, the resonance of the loaded piezo will be shifted due to the extra mass added.

Assuming a damped harmonic oscillator with:

$$\omega_r = \omega_0 \sqrt{1 - 2\zeta^2}, \quad (5.1)$$

where ω_r and ω_0 the resonances of the damped and undamped resonators respectively and ζ the damping ratio. Then, comparing the damped resonance frequencies of the loaded and unloaded piezo, we obtain:

$$\frac{\omega_{r, \text{loaded}}}{\omega_{r, \text{unloaded}}} = \sqrt{\frac{m}{m + M'}} \quad (5.2)$$

where m the mass corresponding to the unloaded piezo resonance and M the mass of the load. Thus, the resonance frequency is expected to decrease. However, the mass of the load is not known to great accuracy. Furthermore, because of potential differences in the exact mounting stiffness of the piezo-mirror system for the different cavities, it is important to determine these resonances experimentally.

Because the contraction of the piezo and the field inside the piezoelectric material are directly related, the impedance will directly reflect actuation of the mirror. To this end, the impedance is measured for a large range of frequencies from a few Hz to a few hundred kHz. The complex impedance is defined as

$$Z = \frac{V}{I}, \quad (5.3)$$

where V and I are the complex voltage and current respectively. Defined as complex quantities in polar notation:

$$V = |V|e^{i(\omega t + \phi)}, I = |I|e^{i(\omega t + \psi)}. \quad (5.4)$$

Because we are only interested in the frequency of the resonance, we use the magnitude of the complex impedance which is simply given by the ratio of the magnitude of the voltage and the current. The impedance is measured using the so-called IV method. This is a simple method where a reference resistance is used over which the voltage drop is measured in tandem with the current [31]. See figure 5.1 for the experimental setup. Using basic circuit theory the impedance is then given by $Z_x = \frac{V_2}{V_1 - V_2} R$. By sweeping the frequency of the AC voltage applied to the piezo and continuously measuring V_1 , V_2 and I , the impedance curve as plotted in figure 5.2 is constructed. For each frequency, the AC voltage is applied for five seconds after which the measurement is performed. This is done to give the system time to settle to a stable motion. The impedance curve shows a strong $1/f$ relation shown by the linear fit with a slope of -0.97 . This is caused by the capacitance of the piezo. A purely capacitive load has an impedance of $Z_C = \frac{1}{j\omega C}$ where j is the imaginary unit, ω the angular frequency and C the capacitance. The fit results in a capacitance at




Figure 5.1: Circuit diagram for IV-method. By measuring the voltage on two points V_1 and V_2 and the current supplied by the power supply the impedance can be calculated. The voltage $V_{ac}(f)$ is swept over various frequencies.

these frequencies of $C_p \approx 388$ nF, which is close to the reported (unloaded) 510 nF in the datasheet. A large resonance at 88 kHz is observed as well as a smaller resonance at 2.5 kHz. At higher frequencies inductive properties dominate as reflected by the increase in impedance. No distinctive features have been discovered beyond the main resonance depicted in this figure and therefore we conclude that the peak at 88 kHz is indeed the main resonance.

These curves have been measured for all eight input- and output mirrors that give similar results. Due to time constraints, these measurements did not use the IV-method but simply the voltage and current directly. It is however to be noted that although many of the piezo-mirror systems show resonant-like features at a few kHz, their exact location can vary from piezo to piezo. Furthermore, some show resonances at tens of kHz, something that is not observed in the first input mirror as in figure 5.2. See figure B.3 for an overview of the impedance curves of all piezos in the filter cavities. These differences are attributed to system-specific mounting stiffness of the piezo-mirror system.

To solidify the link between these impedance spikes and excitations of the piezo, the spectrum of the reflection signal during a lock is compared to the impedance curves. The reflection spectrum during a locked state will reflect vibrations of the piezo as these are directly related near the resonance. The result is plotted in figure 5.3. The locking procedure and how this is achieved is elaborated on in section 5.2.

The top panel shows a large peak at the modulation frequency f_m , which is inherent to dither locking. Its first harmonic at $2f_m$ is also clearly visible. The peak at f_m shows sidebands at 825 Hz and 2.5 kHz at which

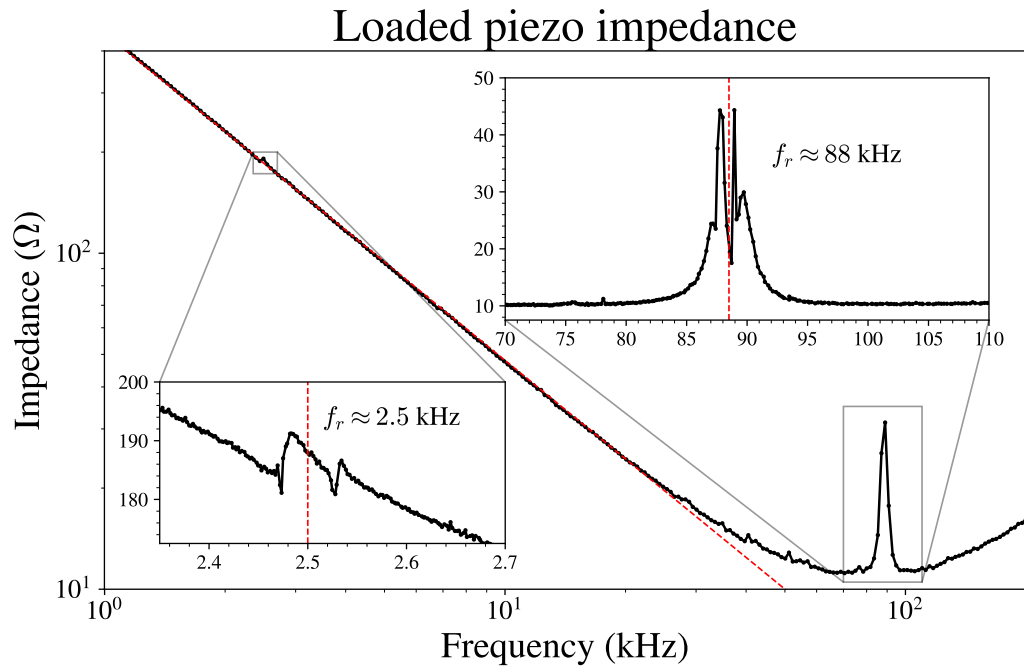


Figure 5.2: Impedance of the piezo attached to the input mirror on the first cavity for a range of frequencies. A straight line is fitted to the impedance curve up to 30 kHz which has a slope of -0.97 showing a strong $1/f$ relation which is typical for the capacitive nature of the piezo at these frequencies. The capacitance of the piezo at these frequencies is $C_p \approx 388$ nF. A clear resonance peak is visible at approximately 88 kHz. A smaller peak is present at approximately 2.5 kHz, as visible in the bottom left inset plot. At higher frequencies no distinctive features have been discovered.

frequencies the impedance show distinct features also for input piezo B. Unlike the 2.5 kHz for input piezo B, the output piezo D shows a smaller resonance nearby at 2.2 kHz. In the sidebands of the reflection spectrum these two are indistinguishable. However, there are resonances for piezo D at tens of kHz that are not observed for piezo B. At these frequencies peaks in the reflection spectrum are observed also. This shows the direct link between the features in the impedance curves and mechanical resonances of the piezo mounting system.

Literature reports similar piezo excitations in an optical cavity locking scheme that impact the quality of the lock by reducing the bandwidth of the feedback [32, 33]. The various resonant frequencies are reported to represent different (bending) modes of the piezo element [34].

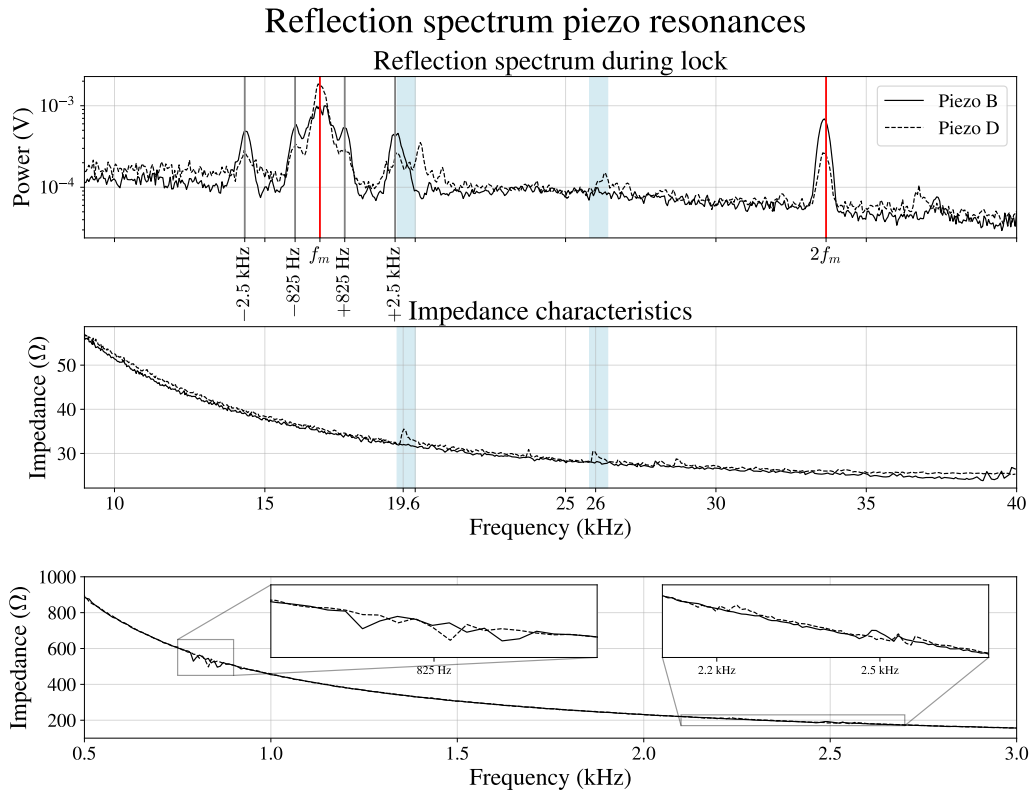


Figure 5.3: A comparison of reflection signal spectra during locking with the impedance curves. The spectra of the reflection signals during locking with input mirror B and output mirror D of the first cavity are in the top plot and the impedance curves of the piezos in the bottom two plots. These measurements are performed with a vacuum pressure below $P \leq 1$ mbar. The modulation at $f_m \approx 16$ kHz (and its harmonics) are clearly visible in the reflection spectrum, indicated with vertical red lines. The peak at f_m shows prominent sidebands at ± 825 Hz and ± 2.5 kHz. At these frequencies the impedance curve in the bottom plot shows features which are not in line with a $1/f$ impedance curve one would expect for a purely capacitive load. More resonances are observed at 19.6 kHz and 26 kHz for piezo D which show peaks in the spectra at those same frequencies. These are not observed in both the impedance curve or the reflection spectrum for piezo B. This indicates that indeed these are likely dynamical features of the loaded piezo.

5.2 Locking procedure

The first step in locking the cavity to the laser source is finding the approximate voltage at which the cavity length is on resonance. To this end, a coarse scan is performed using DAC 2 where its entire range from 0 V to 50 V is swept in 5 seconds. Then, if the reflection signal $r(t)$ is below a threshold R_l , DAC 2 is held at this voltage and the fine scan is initiated. This threshold is typically set slightly above the minimum value of the reflection dip. During the fine scan, DAC 1 is swept from -5 V to 0 V until the resonance is found again. Because the range of the fine scan using DAC 1 is ten times smaller than the coarse scan, DAC 1 is centered at -2.5 V during the coarse scan such that the resonance is always in range of DAC 1. During the fine scan the lockbox is also modulating the piezo to generate an error signal. Note, in order to get the same sign of the phase of the error signal for every lock, an asymmetric scan is performed where the falling edge of the scan is much shorter than the rising edge. Only during the rising edge the lockbox will attempt to lock the cavity. This way, the resonance is approached from the same direction each time and the phase will always have the proper sign. An overview of the procedure to lock the cavity is given in figure 5.4.

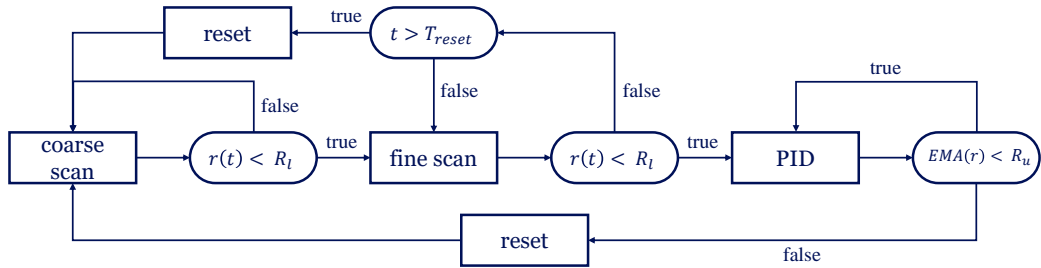


Figure 5.4: A block diagram of the locking procedure. The symbols are defined as follows: $r(t)$ represents the reflection value at time t , R_l and R_u the lock thresholds below/above which the lock is engaged/disengaged and T_{reset} the time in which a resonance should have been detected during the fine scan. First, a coarse scan from 0 V to 50 V is performed until a resonance is found. Then, a fine scan of 0 V to -5 V is performed until a resonance is found. If a resonance is found with the fine scan, the lock is engaged and PID feedback is applied to the piezo. If the fine scan does not find the resonance within T_{reset} , the DACs are reset to the center and the locking procedure starts over again. If the running average $EMA(r)$ reaches the delock threshold R_u , the lock is considered to be lost and the procedure starts over.

When the resonance is found using the fine scan, the lock is engaged and the lockbox starts applying feedback to the piezo. The lock is considered to be lost when the running average $EMA(r)$ of the reflection signal is above the threshold R_u , set typically at 60% of the unlocked (maximum) reflection value. The running average is an infinite response filter of the reflection value known as an exponential moving average:

$$EMA(r)(t) = \alpha \cdot r(t) + (1 - \alpha) \cdot EMA(r)(t - 1), \quad (5.5)$$

where $\alpha \in (0, 1)$ the *smoothing constant* and $r(t)$ the reflection value at time t . The higher the value of α the less the damping is. This filter makes the lock more stable by not considering the cavity to be delocked when sudden noise spikes reach the locking threshold while the PID controller can still correct for it. It is only when the running average reaches the lock threshold that the lockbox will redo the scanning procedure, which generally only happens when the reflection value is at its maximum. Furthermore, by using this simple “expanded” form of the filter, it is very efficient to calculate.

5.2.1 Experimental error signal

An experimentally achieved error signal can be seen in figure 5.5. Here, a sinusoidal modulation with a frequency of $f_{mod} \approx 16.8$ kHz and an amplitude of $100 \mu\text{V}$ is used. For demodulation, a fourth order maximally flat Butterworth low-pass with a cut-off at 400 Hz is employed [35].




Figure 5.5: Experimental error signal during the fine scan. The normalized reflection is plotted in the top plot and the bottom plot shows the error signals after demodulation. The inset in the top plot shows the fine scan resonance without modulation. The in-phase signal is used for feedback. During the fine scan a sinusoidal modulation is applied at a frequency of $f_{mod} \approx 16.8$ kHz with an amplitude of $100 \mu\text{V}$. Here, a fourth order low-pass filter with a cut-off at 400 Hz is used. For the main plots, every tenth point processed by the Teensy is plotted here. The inset plot is measured using an oscilloscope.

The in-phase error signal resembles the (negative) derivative of the reflection signal, which can be used for providing feedback to keep the cavity in lock. A clear linear region in the in-phase error signal is visible, this is where the feedback should be initiated. The oscillations on the reflection dip are not due to electronic noise on the input of the lockbox as they have also been observed on the oscilloscope. Furthermore, as the output of the DACs are quiet as seen in figure 4.5, it is also not likely that these oscillations are caused by noise on the output of the lockbox. These are likely oscillations as the result of excitation of modes in the piezo itself,

which has been elaborated on previously in section 5.1 and is further discussed in section 5.3. The inset plot shows that during a fine scan without modulation, no such strong oscillations are visible, indicating that indeed these are modes of the piezo being excited by the modulation signal.

Note the slight temporal mismatch between the minimum of the reflection dip and the zero of the error signal. This is due to the low-pass filter in the demodulation stage that adds a phase delay to the error signal. This delay should be minimized but is inherent to low-pass filters. However, it is observed that in practice this does not pose a severe problem for locking. This could be explained by the fact that if the lock is engaged outside this linear region while the reflection signal is still around the baseline, made possible by excellent electronic stability of the system, the lockbox is still able to apply feedback to bring the error the setpoint of zero.

A second note must be made on the fact that the quadrature signal is not perfectly zero meaning that not all information is contained in the in-phase error signal as one would expect when considering (3.7). However, in practice, it turns out that the approximation sign in (3.7) is of importance as the filtering is not perfect and some information of higher frequency components are still present in the error signal. This makes the angle optimization imperfect. It has also been observed that despite the asymmetric fine scan and careful consideration of modulation and demodulation waveforms, not all error signals have most of their information contained within the in-phase error signal. This could come from phase mismatches at any point along the feedback loop that cause the phase in the rotation matrix to be off. However, with the right PID parameters, it has been observed that the system is able to lock consistently needing few approaches, indicating that even with a sub-optimal error signal the lockbox is able to lock the cavity to the probe laser.

5.2.2 Locking

Using a modulation frequency of 16.8 kHz and a fourth order Butterworth low-pass filter with a cut-off frequency of 400 Hz, the first filter cavity was successfully locked. The amplitude of modulation was set at 100 μ V. The procedure is plotted in figure 5.6. This lock was achieved while the probe laser laser was locked to the reference cavity with the commercial VU lockbox using the PDH technique where feedback is applied on the piezo actuator of the laser itself rather than on the piezo attached to the the mirror of the cavity. It was found that mostly the integral component is of importance for tuning the parameters of the PID controller.




Figure 5.6: Locking of the first filter cavity. The top plot shows the output voltages of DAC 1 and DAC 2. The bottom figure shows the normalized reflection signal of the cavity. The normalization sets the baseline of the lock at 0% and the reflection out of lock at 100%. The cavity is locked by providing feedback on the input mirror of the cavity. This experiment was performed with a vacuum pressure of 0.3 mbar. Sinusoidal modulation is applied at a frequency of 16.8 kHz and an amplitude of $100 \mu\text{V}$. A fourth order low-pass filter with a cut-off frequency of 400 Hz is used for demodulation. After the coarse scan (DAC 2) finds a resonance at approximately 1 second, it is stopped and the fine scan (DAC 1) is initiated. When the resonance is found with DAC 1 the lock is engaged at approximately 8.5 seconds. Note that this measurement has been performed before the asymmetric scan for DAC 1 was implemented.

Firstly, the coarse scan finds a resonance after which its voltage is fixed. The first resonance reaches only approximately 60% due to the low sampling rate. In reality, this resonance reaches the baseline at 0%, similar to the locked state. Then, after the fine scan finds the resonance, the lock is engaged and the reflection signal drops to a base value. Various noise spikes that are at frequencies of several kHz are visible during the locked state. Environmental acoustic noise, e.g. people talking, can be directly ob-

served in the reflection signal during the locked state. This indicates that acoustic noise is an important factor when considering the quality of the lock. Furthermore, it was observed that at higher vacuum pressures, an increased proportional gain resulted in better locking performance. This highlights the acoustic nature of the noise and shows that this noise also affects the lock at lower frequencies.

Due to the vacuum chamber, the entire optical path is inaccessible which makes it impractical to measure the power transmitted through a single cavity directly. Therefore, as to quantify the quality of the lock, we define the quality of the lock as:

$$Q = 100\% - \frac{\langle r_{lock}(t) \rangle - r_{base}}{\langle r_{unlock}(t) \rangle - r_{base}} = \frac{72.55mV - 65mV}{360mV - 65mV} \approx 97\%, \quad (5.6)$$

where $\langle r_{lock}(t) \rangle$ and $\langle r_{unlock}(t) \rangle$ are the average reflection over 20 seconds during the locked state and unlocked state respectively and r_{base} the baseline of the lock.

It must be noted that to get the actual optical transmission, one should also take into account the coupling of the light to the cavity mode. Furthermore, there are also losses due to optical components that come after the first cavity that will reduce the actual power of the transmitted light incident on the second cavity. These factors make an optical transmission measurement for a single filter cavity difficult.

A better and more interesting measurement would be to measure the transmission through the entire cavity system by measuring the power of the incoming and outgoing light. Such a measurement has been performed and is elaborated on in section 6.2. Considering the context in which this experiment is performed, the quality Q of the lock is important for the subsequent freezing behavior. An ill-behaved lock will not have the benefit of the relaxation of the piezo-mirror system [9]. This will negatively impact the transmission during the freeze lock procedure.

The lockbox has also been tested on the cold cavity that is held at 4 K inside the cryostat. Because of the low temperature and excellent vibration isolation, this experiment is a good benchmark to assess the performance of the lockbox. During such an experiment the probe laser is locked to the cold cavity using the PDH technique, where the error signal is generated by the Zurich Instruments HF2LI lock-in amplifier and fed to the lockbox. Successful locking of the laser to the cold cavity showed that during the unlocked state the reflection signal had an AC-RMS of 8 mV while an AC-RMS of 13 mV was observed during the locked state. With a dark current at approximately 3 mV, this suggests that the lock is close to the limit of the setup noise. This shows the sufficient capabilities of the lockbox.

Lastly, a note must be made on the stability of the closed feedback loop. The Nyquist stability criterion states that the gain at π phase shift should be below unity [36]. This can be intuitively understood from the fact that at π phase shift, the negative feedback is reversed and will make the system unstable if the feedback is allowed to built up above unity gain. However, the loop includes a fourth order low-pass filter that adds significant gain at π phase shift. Experimentally, it was shown that still the system could be locked. Therefore, this indicates sufficient gain margin of the entire closed loop system.

5.3 Feedback bandwidth limitations

Because of the low-pass filter cut-off at 400 Hz, the error signal will not contain features at higher frequencies. This makes the lockbox unable to react to noise with frequencies above this cut-off. Since acoustic noise in the kHz-regime has been observed during locking, a significant amount of effort has been put into increasing the cut-off frequency and attempting to lock the cavity. However, increasing the cut-off also requires an increased modulation frequency for the demodulation process to be effective. But because of the maximum sampling frequency of approximately 1.4 MHz one cannot modulate at any desired frequency. Especially when modulating with a sinusoidal waveform, many points are needed in order to accurately represent the smoothness of a sine without introducing unwanted frequencies. This can be circumvented by modulating with a square wave, which only requires two points to be represented perfectly. This yields a maximum modulation frequency of approximately 700 kHz. Then, by demodulating at any of its (odd) harmonics, an error signal can be retrieved. Although this was experimentally achievable, no successful lock has been realized with either sinusoidal or square wave modulation with high cut-off frequencies of several kHz.

Empirically, the reason for this seems to be the presence of strong 2.5 kHz noise that is visible near the resonance when the feedback is engaged. An example of such an observation can be seen in figure 5.7.

The figure shows an attempted lock with a cut-off at 3 kHz, where clear oscillations at 2.5 kHz are visible as soon as the lock is engaged. This can be explained by the fact that if the bandwidth of the feedback increases beyond this natural minor resonance of the piezo-mirror system, the resonance is excited and the piezo starts to oscillate at this frequency. Because these oscillations are intrinsic to the piezo itself, the feedback loop cannot compensate properly for these oscillations.

Because oscillations in the reflection dip were visible when modulating, the effect of modulation on the excitation of these resonances has also been investigated. This is done by comparing the spectrum of a lock with a strong- and weak sinusoidal modulation and a low cut-off frequency of 400 Hz. The result can be seen in figure 5.8.

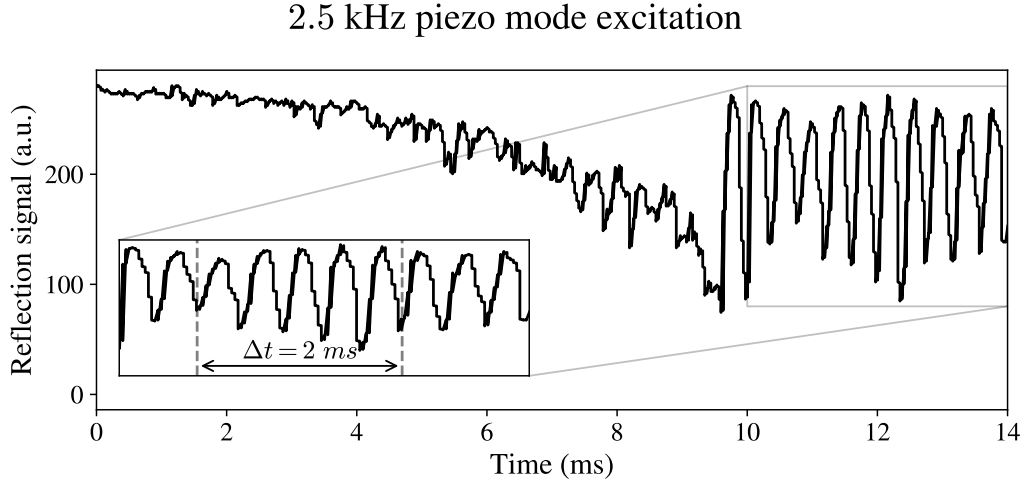


Figure 5.7: Oscillations in the reflection signal at 2.5 kHz. An attempted lock with an increased cut-off frequency of 3 kHz. A clear oscillation at 2.5 kHz is visible as soon as the lock is engaged at around 10 ms. The inset plot shows the oscillations with $\Delta t = 2$ ms for five periods, i.e. 2.5 kHz.

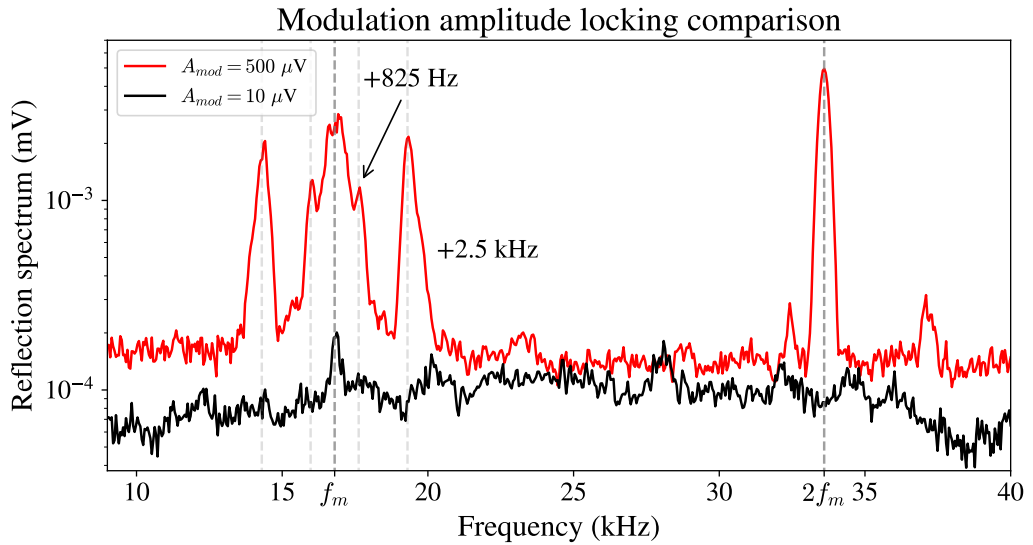


Figure 5.8: Reflection signal spectrum during locked state for strong and weak sinusoidal modulation. The red spectrum is for a modulation amplitude of $A_{mod} = 500 \mu\text{V}$ and the black spectrum for $A_{mod} = 10 \mu\text{V}$. Like in figure 5.3, clear sidebands at ± 2.5 kHz and ± 825 Hz are visible at the modulation frequency $f_m \approx 16.8$ kHz. The sidebands at 825 Hz are much less prominent than the 2.5 kHz sidebands. A large harmonic $2f_m$ is visible. These measurements were performed with a vacuum pressure of $P \approx 3.5 \cdot 10^{-2}$ mbar.

Figure 5.8 shows that a lower modulation amplitude results in less noise in the entire spectrum, which also highlights one of the drawbacks of the dither locking technique. Furthermore, the sidebands at ± 2.5 kHz as well as the strong harmonic at $2f_m$ are significantly lower. The cause of such a large first harmonic at $2f_m$ for this measurement is unknown and has not been observed for other measurements with similar setups. This result indicates that also the modulation signal, which is required for dither locking, is able to excite resonances of the piezo-mirror system even while modulating at frequencies that do not directly correspond to resonances in the impedance curve. This can be explained by the fact that the modulation is not a perfect sine and contains other frequencies at various amplitudes as can be seen in figure 4.4. This result corresponds to the findings as depicted in figure 5.6. Note that the smaller modulation amplitude of $10 \mu\text{V}$ corresponds to an amplitude of only two bits on DAC 1, indicating the excellent noise performance of the lockbox.

The above findings have led to the conclusion that the bandwidth of the feedback loop should be limited as to not excite mechanical resonances in the piezo-mirror system by PID feedback. Because low pass filters always let through some higher frequency components, a cut-off frequency well below 2.5 kHz is required. Furthermore, higher modulation amplitudes resulted in excitation of the resonances despite modulating at a different frequency than the resonance frequency. Therefore, the modulation amplitude should be reduced as much as possible when locking the filter cavities.

Although literature reports successful digital filtering of the feedback at such resonances for a similar system [33], this has been attempted without success. The experiments involved applying a digital notch filter at 2.5 kHz on the output of the lockbox. However, these experiments have not been tried very thoroughly as this was not the main focus of the research at the time. This could therefore still be a promising way to reduce excitation of the resonance at 2.5 kHz and get around the bandwidth limitations.

Chapter 6

Four filter cavities locking

The final results of this work, locking all four filter cavities, are covered in this chapter. The procedure to coordinate the lockboxes and consecutively lock the cavities is explained. The results of an optical transmission measurement while locking all four cavities are discussed. Lastly, results of a lock freeze state for three cavities are discussed.

6.1 Locking procedure

During the time in which the research was conducted, three lockboxes were built since there were only three complete sets of components available at the time. Therefore, the first cavity is locked by the commercial VU lockbox using the PDH technique by locking the probe laser laser to the cavity. The Zurich Instruments HF2LI lock-in amplifier modulates the phase of the light using the Thorlabs EO-PM-NR-C2 EOM at 9.55 MHz. The lock-in amplifier then performs the demodulation using a second order low-pass filter at 12 kHz to generate the error signal. The commercial lockbox provides feedback to the piezo element on the laser crystal. The bandwidth of this piezo element is 100 kHz, much larger than the low-pass corner frequency used in the demodulation by the commercial lockbox. This results in a feedback bandwidth of 12 kHz.

The other three cavities are locked using the dither locking technique performed by the custom built lockboxes. Because of the previously discussed mechanical resonances of the piezos at several kHz, the demodulation low-pass filter of order 4 is reduced to 250 Hz. Although the resulting feedback bandwidth will reduce the quality of the lock as acoustic noise cannot be compensated for, this did result in easy locking of all cavities

and still remain high lock qualities of over 95% for all cavities. Additionally, because it was observed that the modulation signal can excite various modes of the piezo as well, a decreased modulation frequency of 1.9 kHz with an amplitude of 100 μV is used. The same modulation frequency is used for all lockboxes. Although this could theoretically cause interference of the error signals for the consecutive cavities, this was observed to not be a problem for locking. This is likely due to the fact that a small modulation of 100 μV is used, which is only a few percent of the linewidth of the individual filter cavities. Furthermore, all lockboxes use the same (de)locking thresholds and PID parameters. The PID parameters used are: $P = 0.25$, $I = 35$ and $D = 0$. It was found that the derivative term did not improve the locking quality. The integral term has shown to be the most important parameter, similarly to the locking of a single filter cavity. A small proportional component did also improve the lock. The values for these parameters can intuitively be understood from the fact that the entire filter cavity system was designed to have a built-in mechanical low-pass filter and noise conduction through air is reduced by having the system inside a vacuum chamber at $P = 0.6$ mbar [9]. Therefore, the system can be locked using mostly low frequency feedback.

The locking procedure is as follows. First, the probe laser is locked to the first cavity by the commercial VU lockbox. Only when the first cavity is locked, light will be incident on the second cavity. Then, the second lockbox will start the locking procedure as described in 5.2, and so on. The individual custom lockboxes can be started or stopped by sending them commands over HTTP using the custom-built API and user interface. The reflection signals of each cavity during a successful lock of all four cavities can be seen in figure 6.1.

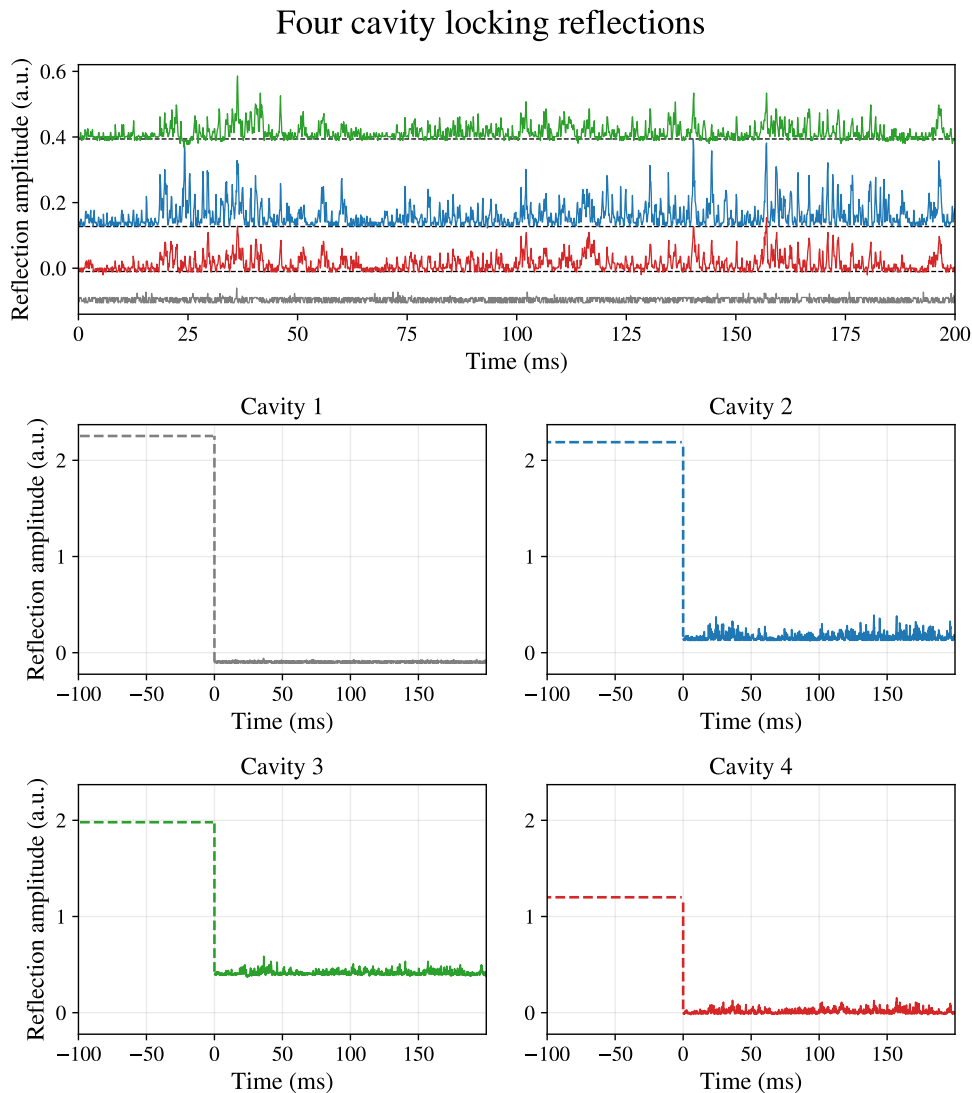


Figure 6.1: Reflection signals and out of lock levels of all four cavities while locked. The top panel shows the measured reflection signals during lock. The reflection of the first cavity is offset by 0.1 in the top plot for visualization purposes. The bottom four plots show the same data but with their out of lock reflection values. The dotted lines represent the out of lock reflection values and the subsequent drop in reflection when locked. The first cavity is held on resonance by locking the laser to the cavity using the PDH method, performed by the commercial VU lockbox. The other three cavities are locked by feedback on the piezo attached to the output mirror of each cavity. Sinusoidal modulation is applied at a frequency of 1.9 kHz and an amplitude of $100 \mu\text{V}$ for the locks of cavity 2 - 4. These measurements were performed with a vacuum pressure of $P \approx 6.5 \cdot 10^{-1}$ mbar. Cavity 2 and 3 show a baseline that is above the dark voltage at approximately 0.

Comparing the locked reflection signal of cavity 1 with the other locked reflection signals, the commercial lockbox achieves great locking performance with the PDH method with a quality of $Q_1 \approx 99.5\%$. Although visibly worse, the locking quality of cavity 2 to 4 are all above 95%.

Both cavity 2 and 3 show a baseline of the lock that is above the dark voltage of the detector at approximately 0. Although this could happen if the setpoint of the feedback loop is not set properly, this has been proven to not be true as a simple scan of cavity 2 and 3 showed that the reflection dip only reaches this same baseline level. This could indicate an optical misalignment in that the field is not perfectly coupled to the mode for cavity 2 and especially cavity 3.

It is interesting to note that many of the noise spikes, as visible in the top plot, are observed in cavities 2 - 4 that are all held on resonance by the custom lockboxes. As these noise spikes are at frequencies of a few kHz, it is likely that these spikes are caused by acoustics that affect all cavities. The first cavity shows a very clean signal with only a few small spikes at e.g. 35 ms as most of the noise can be compensated for by the commercial lockbox due to its much larger bandwidth of 12 kHz. A calculation of the Pearson product-moment correlation coefficients yields the correlation matrix for the four signals:

$$\rho = \begin{bmatrix} 1 & 0.026 & 0.020 & 0.030 \\ 0.026 & 1 & 0.580 & 0.750 \\ 0.020 & 0.580 & 1 & 0.773 \\ 0.030 & 0.750 & 0.773 & 1 \end{bmatrix}, \quad (6.1)$$

where each row/column i corresponds to the reflection of cavity i . The correlations between the signals for cavity 2, 3 and 4 are high, indicating that indeed these spikes likely have a common source. The first reflection signal is much less, if at all, correlated to any of the other reflection signals. This is easily understood when considering the quality of the first lock.

A second interesting note can be made. When inspecting the reflection of the second cavity, there are spikes in the signal that do not have an accompanying spike in the reflection signals for cavity 3 and 4, e.g. at $t \approx 25$ ms. This can be understood by the fact that if the lock is slightly off for cavity i , there will be less light incident on cavity $i + 1$, reducing the reflected signal for cavity $i + 1$. Such spikes are thus not caused by common noise but rather by inaccuracies in the lock of that specific cavity. There are few instances of such spikes, suggesting minimal idiosyncratic noise sources, e.g. electronic noise.

6.2 Optical transmission

The figure of merit for this work is the optical transmission through the entire filter cavity system. The power going into the fiber, that is the input to the filter cavity system, is measured to be $P_{in} \approx 42.0 \mu\text{W}$. During the locked state, the transmitted power is measured to be $P_{out} \approx 16.8 \mu\text{W}$. The dark current was measured to be $P_{dark} \approx 0.5 \mu\text{W}$. This results in an optical transmission of $T \approx 40\%$.

As discussed in section 6.1 and evident from figure 6.1, the quality of the locks is high, meaning that most power is likely lost due to optical misalignment. Although optical absorption and scattering on the many optical surfaces in the filter cavity system could also introduce losses, these contributions are considered to be minimal because of the high quality components that were shown to have only 4% loss through the entire system of optical isolators [9]. The high baselines of the locks for cavity 2 and especially cavity 3 indicate an optical misalignment, as can also be seen by the lower out of lock value of cavity 3 and 4. The out of lock value for cavity 4 is approximately 39% lower than for cavity 3, which corresponds to an optical attenuation of approximately the same percentage as identical detectors, transimpedance amplifiers and bias voltages are used. Furthermore, the power that goes into the filter cavity system is inevitably measured before a E2000 fiber to fiber female adapter that have a typical loss of a few percent [37]. These observed factors suggest that optical misalignment is the main bottleneck for optical transmission through the filter cavity system.

Research done by I. Galinskiy et al. (2020) on a very similar setup with four cascaded optical cavities showed 30% transmission, while having a broader linewidth at approximately 30 kHz [38]. They reported similar findings in that the optical transmission is mostly limited by optical alignment. Comparing the transmission figures, this shows excellent performance of the system.

6.3 Lock freeze

As explained before, during the collection of Stokes photons, the probe laser must be turned off and the lock of the filter cavities must be frozen. During this freeze state, the lockboxes do not receive any reflection signal and thus no active feedback can be applied. As a first experiment, the freeze state is implemented by leaving the voltages applied to the piezo constant during the freeze. Figure 6.2 shows the normalized reflection

signals during a freezing experiment in which the first three cavities are frozen. The probe laser was locked to the reference cavity during the entire experiment.

The transmission retains over 90% of the transmission during the locked state for a period of 4.9 seconds. Although during single photon experiments all four cavities need to be frozen instead of three, this does give an initial indication of the cavity stability and timescales that can be expected during the frozen state. As STIRAP experiments have a thermal decoherence time at 1 K on the order of milliseconds, such freeze lock performance is satisfactory [9].

As evident, the main contribution for the loss in transmission over time is the third cavity. The first two cavities remain close to the resonance for the entire period of 30 seconds, losing only 20% on the first cavity and 3% on the second cavity. One of the causes for the higher drift in the third cavity could be the fact that this cavity has been locked for the least amount of time. It has been shown that locking the cavities for an extended period causes the feedback loop to reach a steady state, possibly due to relaxation of mechanical vibrations of the piezo [9]. To minimize this effect, the same experiment should be repeated but instead with much longer locking times. The difference in drift between the three cavities could also be caused simply by the stochastic nature of this process.

It is interesting to note that the drift of the resonance is unidirectional. This could indicate that the drift is caused by e.g. a thermal gradient in the lab or piezo relaxation, a common phenomena for piezos in which the dipoles in the piezoelectric relax, causing the piezo to lengthen over time [39]. This suggests that more advanced freezing behavior, e.g. linear extrapolation of the applied feedback, could potentially increase the transmission curves during the frozen state.

Finally, it is important to note that this measurement has a stochastic nature. Repeated measurements of such transmission curves could give varying results due to environmental fluctuations. The resulting distribution of transmission curves will give a deeper understanding of the freezing performance.




Figure 6.2: Transmission through the first three cavities during the frozen state. The top plot shows the normalized transmission of three cavities, smoothed with a Savitzky Golay filter with window size 100 datapoints and a polynomial order of 2. The normalized reflection R_i of cavity i is in the bottom plot. The normalization includes a correction for the baseline reflection during the locked state. The piezos are frozen on $t = 0$ and a constant voltage is applied. All three cavities were locked for approximately 1 minute before the locks were frozen. The laser used was locked to the reference cavity during the entire experiment. These measurements were performed with a vacuum pressure of $P \approx 4.0 \cdot 10^{-1}$ mbar. During a period of 4.9 seconds, the transmission is over 90% of the transmission during the locked state.

Discussion

7.1 Recommendations

Various electromechanical resonances of the piezos have been observed through measuring the impedance for a large range of frequencies. However, as most of the time has been spent on investigating the resonances of piezos of the first cavity, it could be worthwhile to measure the impedance curves as in figure B.3 at a higher resolution. Then, if for a piezo no resonances are found at low frequencies, the bandwidth of feedback for that cavity can be increased. This could result in a higher lock quality.

The four cavity locking experiment was performed at a relatively high pressure of $P \approx 6.5 \cdot 10^{-1}$ mbar. This was the result of a small leak in the fiber feed-through on the top of the vacuum chamber. Such high pressures are observed to result in more acoustic noise by conduction through air. Since pressures of 10^{-4} mbar are achievable in the vacuum chamber, the vacuum could be improved to reduce the noise on the locks.

To further increase the quality of the lock, a higher bandwidth feedback loop could be realized by leveraging digital or analog filters to avoid excitation of the resonances. At the time of writing, the first successful locks employing digital notch filters and an increased feedback bandwidth of 3 kHz are being realized. In this experiment, the notch filters are placed at the frequencies where peaks are observed in the spectrum of the reflection signal during locking. The initial results show a higher quality of the lock. More experimentation is needed to reach strong conclusions on the working of such notch filters. Additionally, literature reports successful implementation of a custom design digital filter where poles and zeroes are placed such that its transfer function closely resembles the inverse of the response of the system [33]. This could be an interesting method to

prevent spurious excitation of the resonances, allowing for an increased feedback bandwidth.

Instead of bandwidth increases to compensate for acoustic noise, another way to improve the quality of the lock is to reduce the noise itself. This can be done by using a higher quality vacuum, removing noise conduction through air. As the current filter cavity system is connected directly on the optical table using stiff materials, noise can be conducted through the vacuum chamber material itself. Mechanical damping through the use of springs can be used to reduce this [38]. Lastly, as the the system is in a noisy lab room without any acoustic shielding, extra attention could be given to soundproofing the filter cavities against environmental noise.

The reflection signals during the locking of all four cavities showed that optical misalignment is a large component of the optical loss through the cavities. Careful realignment of all the optical components throughout the entire filter cavity system, especially for cavity 2 and 3, could yield several tens of % added optical transmission.

Furthermore, several parameters of the individual lockboxes can be optimized to achieve higher quality locks. Although it was shown that all cavities can be locked using the same modulation frequency, using different modulation frequencies could result in even higher quality locks. Moreover, the low-pass filter is placed at 250 Hz whereas higher cut-offs could still result in successful locking while increasing the bandwidth. The author is convinced that optimization of the these parameters could improve the quality of the locks.

Due to time restrictions on this research, the transmitted signal has not been measured using a photodetector. This signal is interesting to measure for analyzing the performance of the locking as the amount of noise on the transmitted signal will be a direct measure thereof. To this end, an extra detector should be aligned with the output fiber of the filter cavity system.

In order to allow for a full lock freeze operation measurement, a fourth custom lockbox needs to be built. Then, similarly to the experiment performed in this work, the transmission curve through the entire system can be measured. Due to time restrictions of this research, only a single transmission curve was measured. To gain deeper insight into the statistical nature of the length stability of the individual cavities, repeated measurements of such transmission curves should be conducted. The resulting distribution can give a better understanding of the freezing behavior. Furthermore, as the length drift of the cavities has been observed to be unidirectional, more advanced lock freeze behavior could improve the transmission curves. An obvious, simple and likely effective first attempt would be to linearly extrapolate the applied feedback of the locked state.

7.2 Conclusions

By careful consideration of electrical components and the grounding thereof, a low-noise lockbox has been developed with a resolution of $5\ \mu\text{V}$ and a noise floor of $3\ \mu\text{V}$. Important features like the input amplification, a custom I/O extension board, two Teensys and HTTP communication allow for flexible operation and coordination of the lockboxes.

The various electromechanical resonances of the piezos have proven to limit the locking capabilities of the lockbox. By reduction of the feedback bandwidth we demonstrated the locking of a single filter cavity and achieve a high quality lock.

Using a reduced bandwidth of only 250 Hz and a low modulation frequency of 1.9 kHz all four cavities were locked with a high lock quality, resulting in an optical transmission of 40%. The reflection signals of the individual filter cavities showed that the optical transmission seems to be mainly limited by optical alignment. The first lock freeze operation showed that during a period of 4.9 seconds, 90% of the original transmission can be retained through three filter cavities. This leaves sufficient room for STIRAP experiments that have typical timescales on the order of milliseconds.

Summary

This work contributes to literature by analyzing peaks in the impedance curves of the piezos in the filter cavities and attributing them to electromechanical resonances of the piezo-mirror system. By comparison of the reflection signal spectrum during the locked state with the impedance curves for two piezo-mirror systems, the connection between the aforementioned peaks and dynamical resonances is solidified. Furthermore, a detailed description of the design of a custom built lockbox with an output resolution of $5\ \mu\text{V}$ is given. It was shown that this design yields a low noise floor below the least significant bit of the 20 bit DAC. Using this design a high quality lock was realized. Important locking capability limitations caused by the aforementioned resonances are noted. To avoid excitation of such resonances, a reduced bandwidth of 250 Hz and a modulation frequency of 1.9 kHz were used. This allowed for locking all four cavities, achieving an optical transmission of 40%. An initial lock freeze procedure of three cavities showed that 90% of the initial transmission during the locked state can be retained for a period of 4.9 seconds. This is sufficient for STIRAP experiments that have typical timescales on the order of milliseconds.

Acknowledgements

During the course of six months, working with Enrique Morell has taught me many things about electronics, PCB design, programming in C++/C# and feedback loops. I have also learned many things about optics, optomechanical experiments and doing research in general. For this I want to thank Enrique Morell, Leon Raabe and Xinrui Wei. Moreover, I want to express my gratitude towards Prof.Dr. D. Bouwmeester for the many meetings and lab visits. These moments have been instrumental in maintaining a clear perspective and research goal. I am very grateful for the opportunity that was given to me to do research in the Bouwmeester lab. I also want to thank Geert Timmerman for the occasional fruitful discussions during the bike rides home or at the diner table that have helped me formulate ideas more clearly. Furthermore, Ing. K Heeck has been a source of wisdom and experience for many problems that were faced. I want to thank the people from the Electronics Department and the Fine Mechanical Department who have greatly helped pushing this project forward at a high pace by providing invaluable insights and helping with various difficult practical tasks. Both of my brothers Lou and Bram and my girlfriend Sharon have helped by providing valuable feedback on figures and the thesis itself. Lastly, I want to thank my parents for supporting me throughout the duration of this research.

Bibliography

- [1] B. Hensen, H. Bernien, A. E. Dreau, A. Reiserer, N. Kalb, M. S. Blok, J. Ruitenber, R. F. L. Vermeulen, R. N. Schouten, C. Abellán, W. Amaya, V. Pruneri, M. W. Mitchell, M. Markham, D. J. Twitchen, D. Elkouss, S. Wehner, T. H. Taminiau, and R. Hanson, *Loophole-free Bell inequality violation using electron spins separated by 1.3 kilometres*, Nature **526**, 682 (2015), Publisher: Nature Publishing Group.
- [2] T. E. H. T. Collaboration et al., *First M87 Event Horizon Telescope Results. I. The Shadow of the Supermassive Black Hole*, The Astrophysical Journal Letters **875**, L1 (2019), Publisher: The American Astronomical Society.
- [3] LIGO Scientific Collaboration and Virgo Collaboration et al., *Observation of Gravitational Waves from a Binary Black Hole Merger*, Physical Review Letters **116**, 061102 (2016), Publisher: American Physical Society.
- [4] A. Peres and D. R. Terno, *Quantum information and relativity theory*, Reviews of Modern Physics **76**, 93 (2004), Publisher: American Physical Society.
- [5] C. J. Isham, *Canonical Quantum Gravity and the Problem of Time*, in *Integrable Systems, Quantum Groups, and Quantum Field Theories*, edited by L. A. Ibort and M. A. Rodríguez, NATO ASI Series, pages 157–287, Springer Netherlands, Dordrecht, 1993.
- [6] W. Marshall, C. Simon, R. Penrose, and D. Bouwmeester, *Towards Quantum Superpositions of a Mirror*, Physical Review Letters **91**, 130401 (2003).

- [7] C. Pfister, J. Kaniewski, M. Tomamichel, A. Mantri, R. Schmucker, N. McMahon, G. Milburn, and S. Wehner, *A universal test for gravitational decoherence*, *Nature Communications* **7**, 13022 (2016), Publisher: Nature Publishing Group.
- [8] R. Penrose, *On Gravity's role in Quantum State Reduction*, *General Relativity and Gravitation* **28**, 581 (1996).
- [9] V. Fedoseev, *Stimulated raman adiabatic passage in optomechanics*, PhD thesis, Leiden University, 2022, ISBN: 9789085935261.
- [10] M. Aspelmeyer, T. J. Kippenberg, and F. Marquardt, *Cavity optomechanics*, *Reviews of Modern Physics* **86**, 1391 (2014), Publisher: American Physical Society.
- [11] B. ten Haaf, *Filter cavity design for filtering photons at 1MHz separation*, (2021).
- [12] Coherent, *Mephisto/Mephisto S*, 2024.
- [13] N. Ismail, C. C. Kores, D. Geskus, and M. Pollnau, *Fabry-Perot resonator: spectral line shapes, generic and related Airy distributions, linewidths, finesses, and performance at low or frequency-dependent reflectivity*, *Optics Express* **24**, 16366 (2016).
- [14] B. Chris, *RF Circuit Design, Second Edition by Chris Bowick, John blyler and Cheryl Ajluni- By*.
- [15] O. Huisman, *An Exploration into Digital Dither Locking*, (2021).
- [16] R. W. P. Drever, J. L. Hall, F. V. Kowalski, J. Hough, G. M. Ford, A. J. Munley, and H. Ward, *Laser phase and frequency stabilization using an optical resonator*, *Applied Physics B* **31**, 97 (1983).
- [17] E. D. Black, *An introduction to Pound Drever Hall laser frequency stabilization*, *American Journal of Physics* **69**, 79 (2001).
- [18] A. V. Oppenheim, R. W. Schafer, and R. W. Schafer, *Discrete-time signal processing*, Always learning, Pearson, Harlow, third edition, pearson new international edition edition, 2014.
- [19] *AD5252 Datasheet and Product Info | Analog Devices*.
- [20] C. Shannon, *Communication in the Presence of Noise*, *Proceedings of the IRE* **37**, 10 (1949), Conference Name: Proceedings of the IRE.

-
- [21] AD5791 Datasheet and Product Info | Analog Devices.
- [22] D. Krakauer, *Anatomy of a Digital Isolator* | Analog Devices, 2011.
- [23] P. Wilson, *Chapter 6 - Digital Circuits*, in *The Circuit Designer's Companion (Fourth Edition)*, edited by P. Wilson, pages 259–320, Newnes, 2017.
- [24] C. PCB Solutions, *Decoupling Capacitor Placement in PCB Layout*, 2020.
- [25] D.-K. Electronics, *OPA182IDT Zero-Drift Amplifier 1 Circuit Rail-to-Rail 8-SOIC from Texas Instruments*.
- [26] EVAL-AD5791 Evaluation Board | Analog Devices.
- [27] ADA4522-1 Datasheet and Product Info | Analog Devices.
- [28] IEEE Standards Association IEE 802.3, 2018.
- [29] W. M. van Spengen, *The electromechanical damping of piezo actuator resonances: Theory and practice*, *Sensors and Actuators A: Physical* **333**, 113300 (2022).
- [30] Noliac, *NAC2124 Ring Actuator*.
- [31] X. Zhu, L. Zhang, Y. Guo, P. Wei, R. Zai, J. Drewniak, and D. Pommerenke, *I-V Method Based PDN Impedance Measurement Technique and Associated Probe Design*, in *2019 IEEE International Symposium on Electromagnetic Compatibility, Signal & Power Integrity (EMC+SIPI)*, pages 30–33, New Orleans, LA, USA, 2019, IEEE.
- [32] M. Pallmann, T. Eichhorn, J. Benedikter, B. Casabone, T. Hummer, and D. Hunger, *A highly stable and fully tunable open microcavity platform at cryogenic temperatures*, *APL Photonics* **8**, 046107 (2023).
- [33] T. Ruelle, D. Jaeger, F. Fogliano, F. Braakman, and M. Poggio, *A tunable fiber Fabry Perot cavity for hybrid optomechanics stabilized at 4 K*, *Review of Scientific Instruments* **93**, 095003 (2022).
- [34] M. Rost, G. Baarle, A. Katan, W. M. Van Spengen, P. Schakel, W. Loo, T. Oosterkamp, and J. Frenken, *Video-rate scanning probe control challenges: Setting the stage for a microscopy revolution*, *Asian Journal of Control* **11**, 110 (2009).
- [35] S. Butterworth, *On the Theory of Filter Amplifiers*, **7**, 536 (1930).
-

- [36] H. Nyquist, *Regeneration theory*, **11**, 126 (1932).
- [37] Diamond, *Diamond E-2000 Fiber Family*.
- [38] I. Galinskiy, Y. Tsaturyan, M. Parniak, and E. S. Polzik, *Phonon counting thermometry of an ultracoherent membrane resonator near its motional ground state*, *Optica* **7**, 718 (2020), Publisher: Optica Publishing Group.
- [39] R. Changhai and S. Lining, *Hysteresis and creep compensation for piezoelectric actuator in open-loop operation*, *Sensors and Actuators A: Physical* **122**, 124 (2005).
- [40] M. J. Lawrence, B. Willke, M. E. Husman, E. K. Gustafson, and R. L. Byer, *Dynamic response of a Fabry Perot interferometer*, *Journal of the Optical Society of America B* **16**, 523 (1999).
- [41] Y. Dumeige, S. Trebaol, L. Ghisa, T. K. N. Nguyen, H. Tavernier, and P. Feron, *Determination of coupling regime of high-Q resonators and optical gain of highly selective amplifiers*, *JOSA B* **25**, 2073 (2008), Publisher: Optica Publishing Group.
- [42] A. E. Shitikov, I. A. Bilenko, N. M. Kondratiev, V. E. Lobanov, A. Markosyan, and M. L. Gorodetsky, *Billion Q-factor in silicon WGM resonators*, *Optica* **5**, 1525 (2018).
- [43] H. Tavernier, P. Salzenstein, K. Volyanskiy, Y. K. Chembo, and L. Larger, *Magnesium Fluoride Whispering Gallery Mode Disk-Resonators for Microwave Photonics Applications*, *IEEE Photonics Technology Letters* , 5570900 (2010).

Appendices

Cavity ringdown

An interesting and commonly used method to determine the quality factor of an optical resonator is the so-called cavity ringdown method. In this method, the interference between built-up light inside the cavity and reflected light is measured [40, 41]. As the field inside the cavity decreases exponentially, damped oscillations are observed. From the standard damped harmonic oscillator model, one can then easily determine the Q -factor of the optical resonator using:

$$Q = \frac{\omega_0 \tau}{2}, \quad (\text{A.1})$$

where ω_0 the resonant frequency and τ the ringdown time [42] [43]. The ringdown time is defined as the time in which a fraction $1/e$ of the field remains in the cavity.

The above method is realized by quickly scanning over the resonance by sweeping DAC 2 over its entire range (50 V) at 4 Hz. By measuring the reflection signal, the previously mentioned damped oscillation is observed as can be seen in figure A.1.

The clear linear relation demonstrates the expected exponential decay of the absolute amplitude of the peaks. Utilizing a linear fit, a ringdown time of approximately $\tau \approx 13 \mu\text{s}$ is determined, resulting in a measured quality factor of $Q \approx 1.2 \times 10^{10}$. A different definition for the Q -factor is:

$$Q = \frac{\nu_0}{\nu_{FWHM}}, \quad (\text{A.2})$$

where ν_0 is the resonant frequency and ν_{FWHM} the linewidth of the resonance. For $\lambda_0 = 1064 \text{ nm}$ and $\nu_{FWHM} \approx 22.4 \text{ kHz}$, this yields $Q \approx 1.3 \cdot 10^{10}$. This is reasonably close to the ringdown method, showing the correctness and validity of the method.

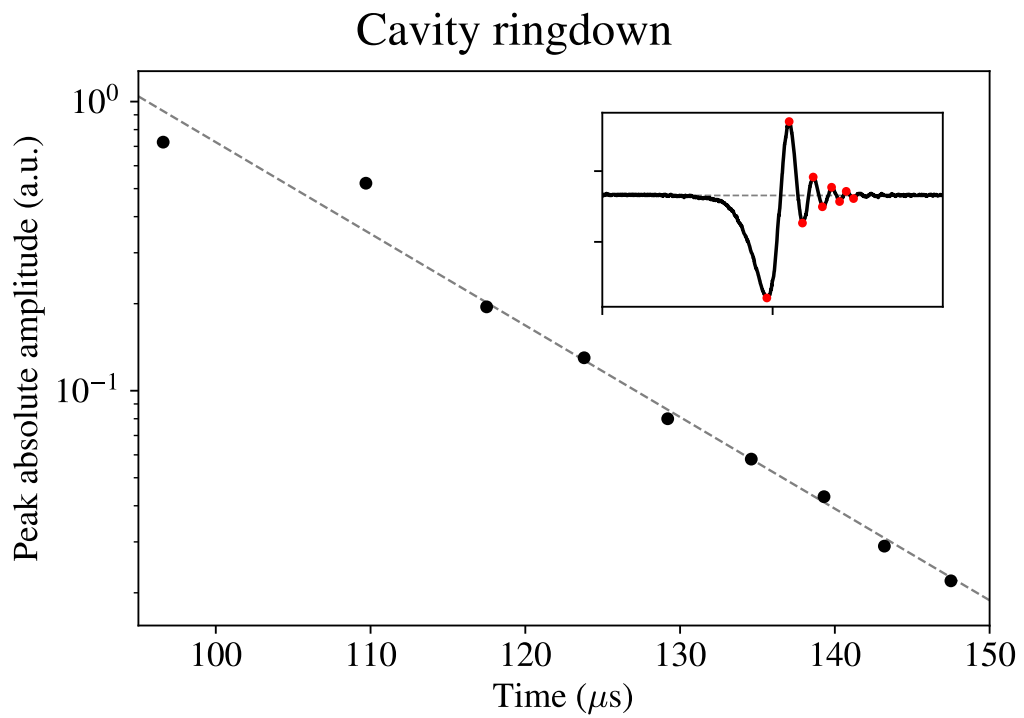


Figure A.1: Filter cavity ringdown measurement by fast piezo sweep. A clear exponential decay of the absolute amplitude of the peaks in the reflection signal as depicted in the inset plot. The fit results in a cavity ringdown time of $\tau \approx 13 \mu\text{s}$. This gives a measured quality factor of $Q \approx 1.2 \cdot 10^{10}$.

Appendix B

Auxiliary figures




Figure B.1: Photograph of the filter cavity system. The entire vacuum chamber is approximately 2 meters long with each cavity being 40 cm. The assembly on top contains LEMO ports for the piezos and detectors inside the vacuum chamber. A vacuum gauge is placed on the other end of the top inlet. The valve on the front allows for pumping the vacuum chamber.



Figure B.2: Photograph of the CF1400 cryostat from Leiden Cryogenics in the Bouwmeester/Oosterkamp lab. The horizontal plates from top to bottom are at increasingly low temperatures starting from 50 K to 10 mK at the bottom of the cryostat. The cavity with the membrane-in-the-middle setup is attached to one of the Eddy current vibration isolators at the bottom. The mK temperatures are achieved using $^3\text{He}/^4\text{He}$ dilution.

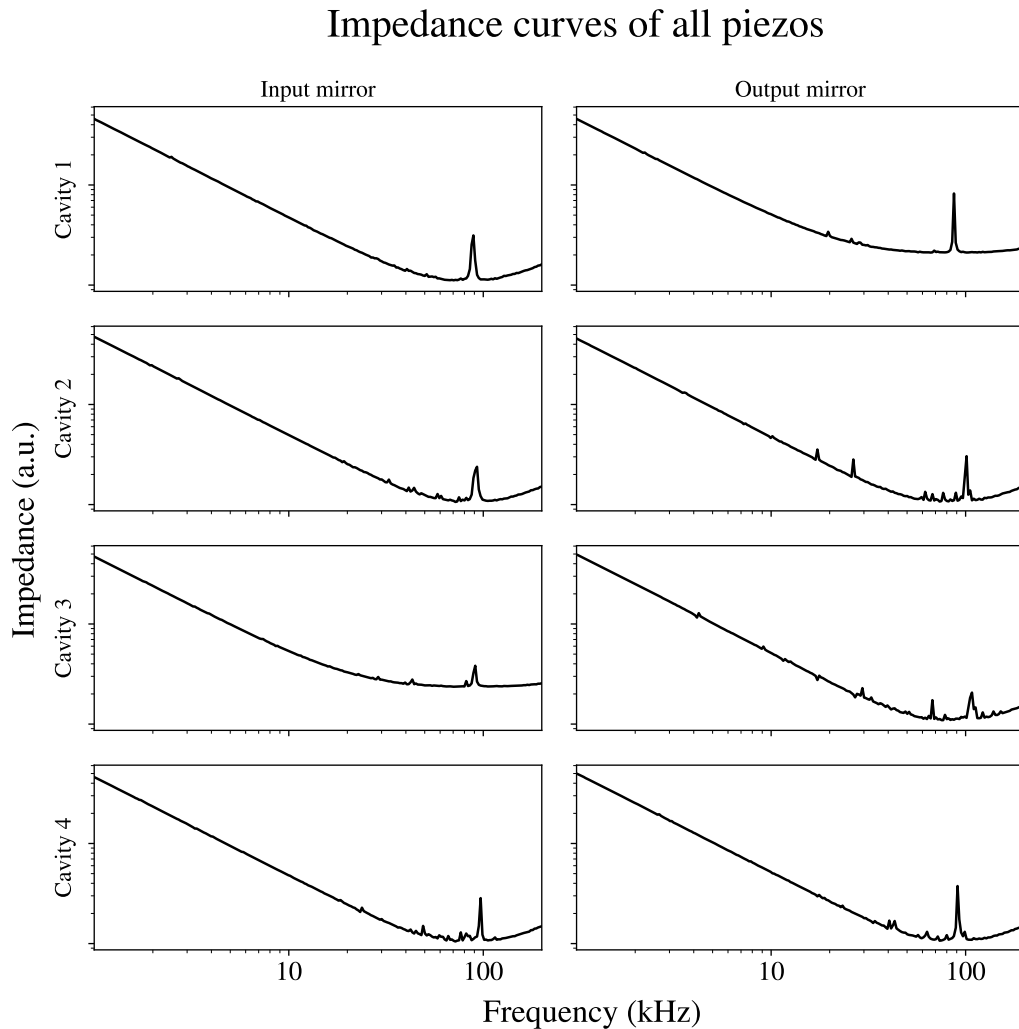


Figure B.3: Impedance curves for the piezos on all piezo-mirror systems. Each row represents a cavity which has both an input and output mirror in the first and second column respectively. The horizontal axis represents the frequency of the AC voltage applied to the piezo and the vertical axis the impedance. Although all piezos show a similar main resonance at approximately 85 - 105 kHz, smaller resonances can vary from piezo to piezo, e.g. the piezo of the output mirror of the second cavity contains resonances at approximately 10 and 20 kHz while other piezos do not show such features. This could be due to differences in exact mounting stiffness of the piezo to the mirror.